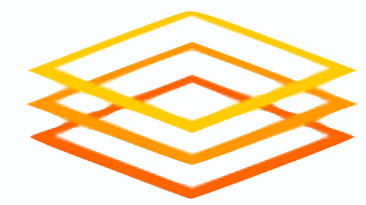


Distributed High Throughput Computing from the Campus Perspective

April 8, 2014

Lothar Bauerdick
OSG Executive Director



**Or, me attempting to
overcome LHC and Labs myopia
and trying to understand Campuses
and Campus Researchers... :-)**

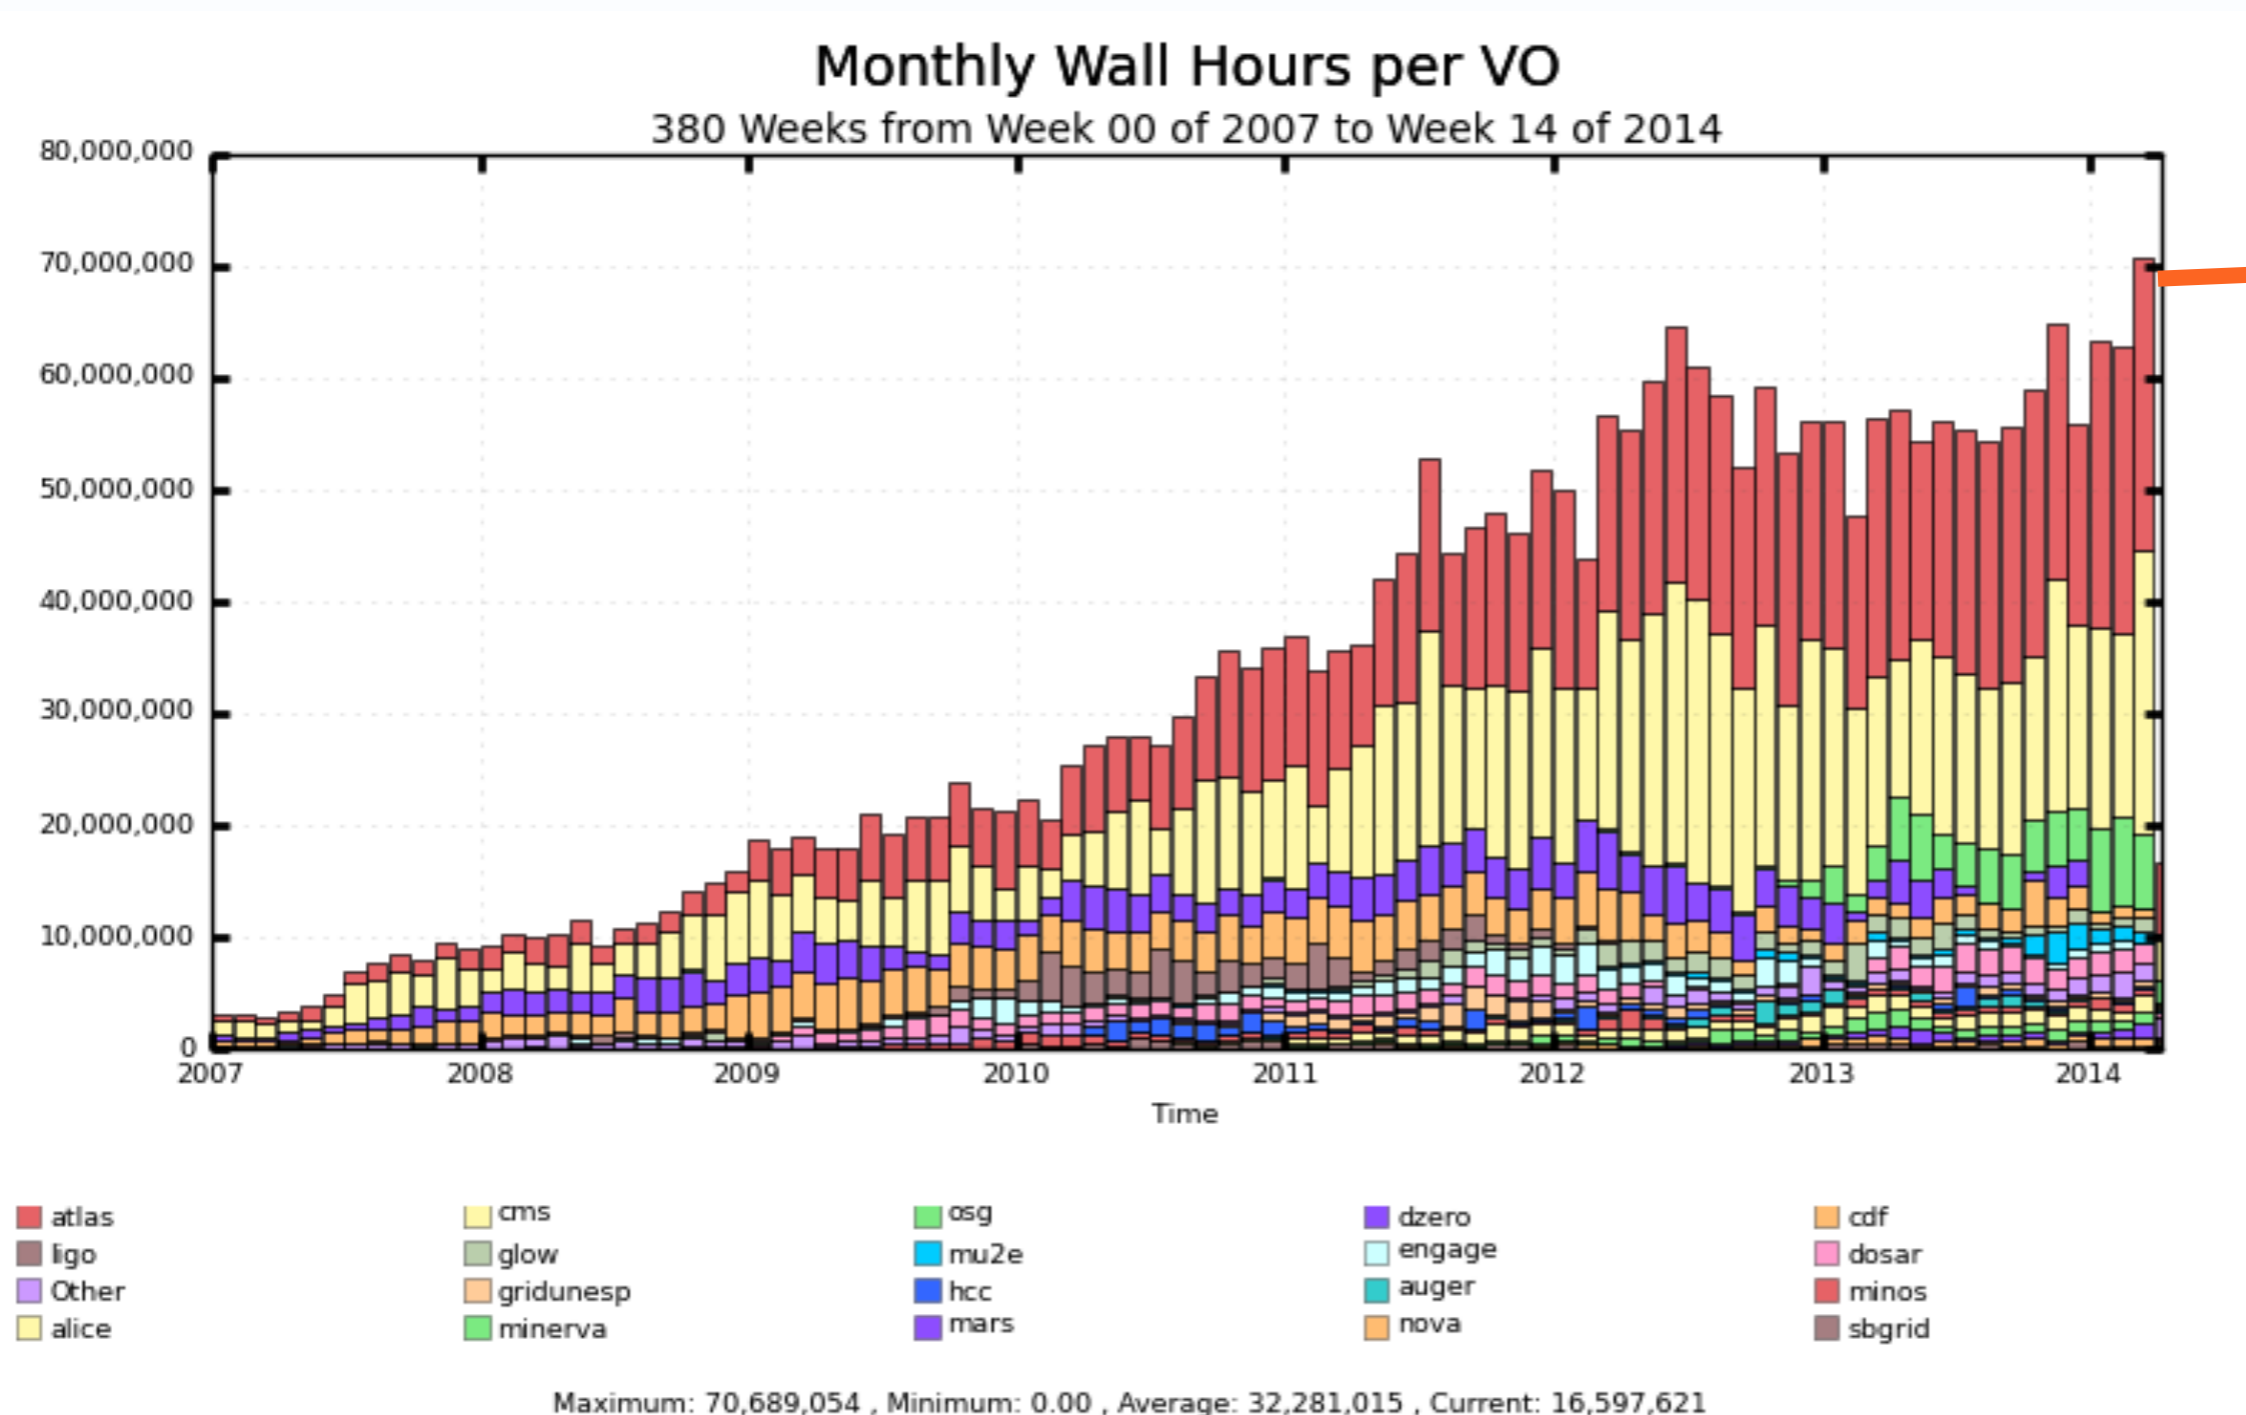


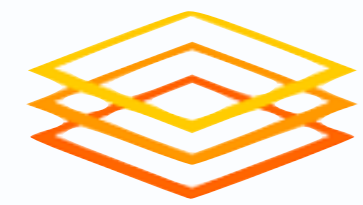
OSG “Size” in Terms of CPU Resources

Open Science Grid

★ LHC plans to ~double CPU resources over coming 2 years

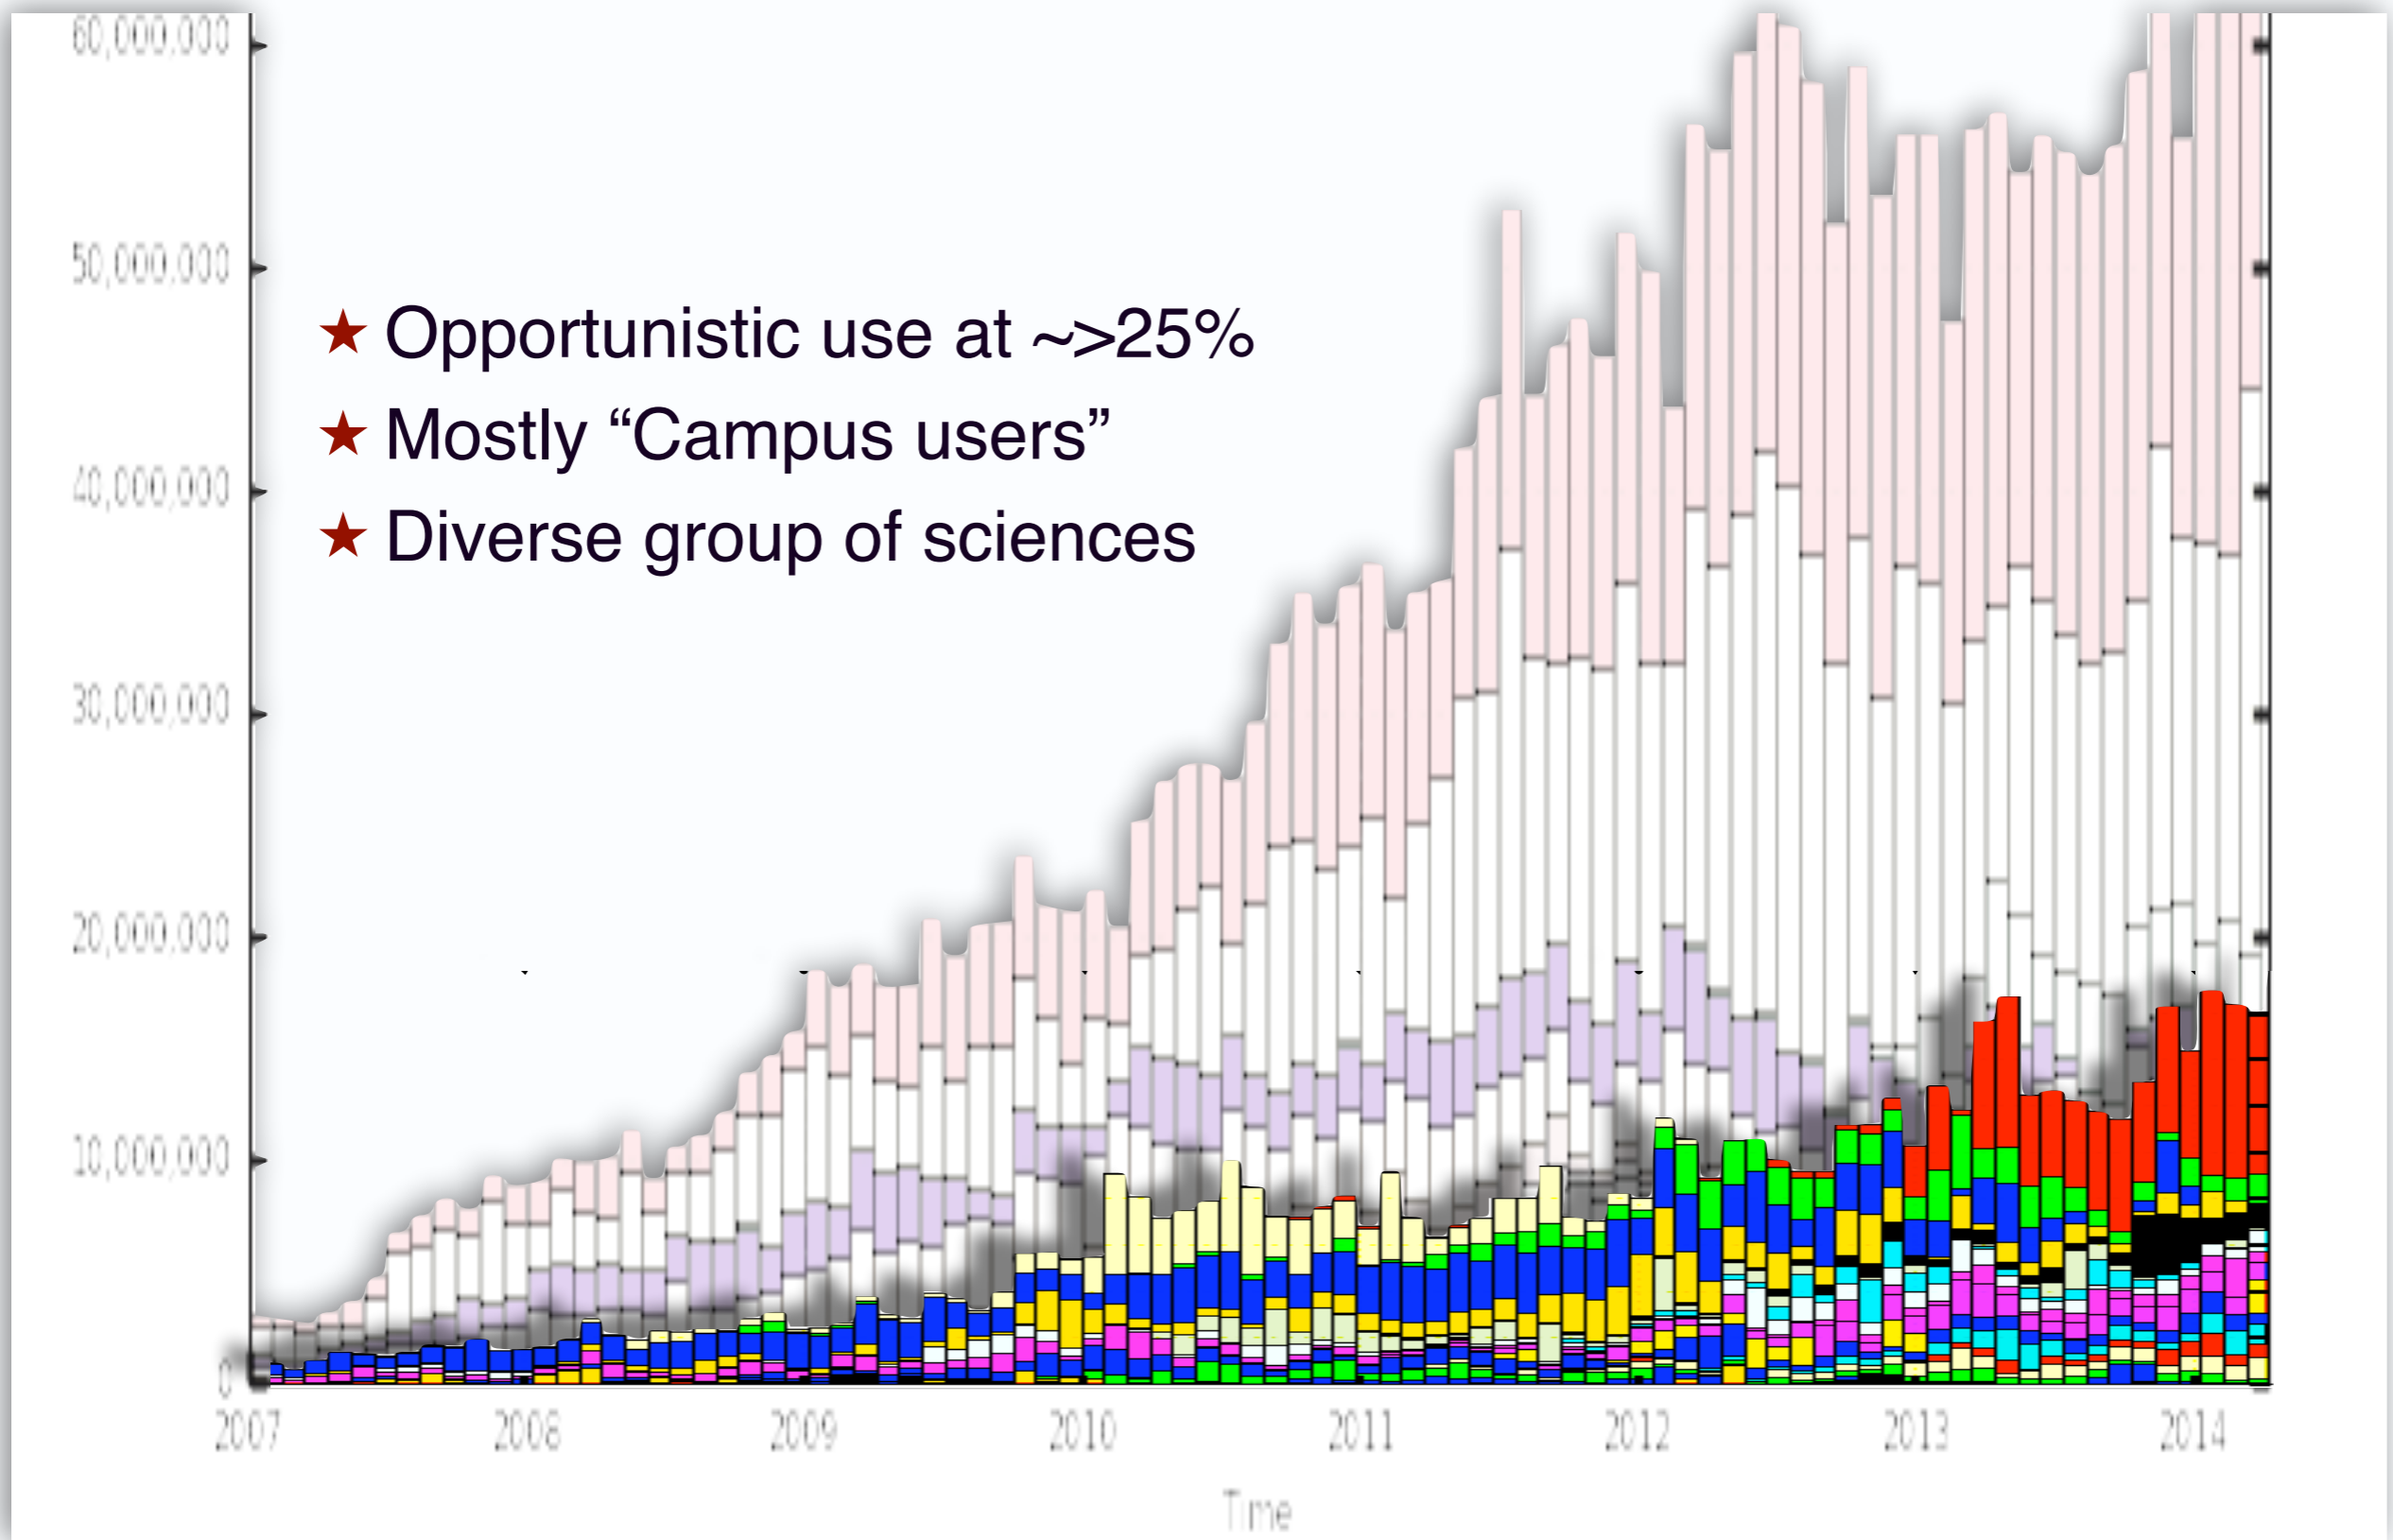
LHC Run2

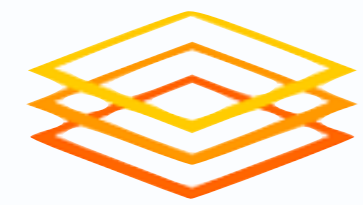




~Proportional Increase of Opportunistic Use

- ★ Opportunistic use at $\sim >25\%$
- ★ Mostly “Campus users”
- ★ Diverse group of sciences

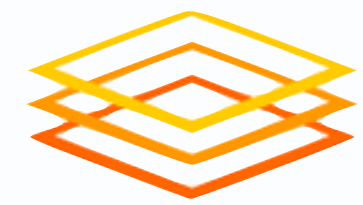




OSG for the "Long Tail of Science"

Open Science Grid

- ◆ "long tail" of researchers that do not have access to sufficient in-house computing resources and/or skills, or that profit from "more"
- ◆ From the 2012 document OSG Future CS Research Needs
 - ★ "OSG will be considered successful if we not only keep up with the growing LHC needs but also expand to enhance the computing throughput to a broad spectrum of scientists at a variety of scales, from individual users at a single campus to multi-institutional experiments
 - ★ Among our scientific user community, we see two extremes of such challenges.
 - ◆ On one extreme, there are the large collaborations in HEP that need innovation to help the transition from petabytes to exabytes of data volumes, and from 10,000 cores to 100M cores within the global DHTC systems over the next ten years.
 - ◆ On the other extreme, there are an ever increasing number and diversity of small to mid-size collaborations of domain scientists that struggle with crossing order of magnitude boundaries in their computing — transitioning from gigabytes to terabytes of data volumes at scales of 1-10 million CPU hours.
 - ★ Over the next 10 years, we expect the number and diversity of groups at this end of the spectrum to increase as data and compute intensive science becomes the norm rather than the exception."



Survey of Who Provides Resources that Researchers Utilize

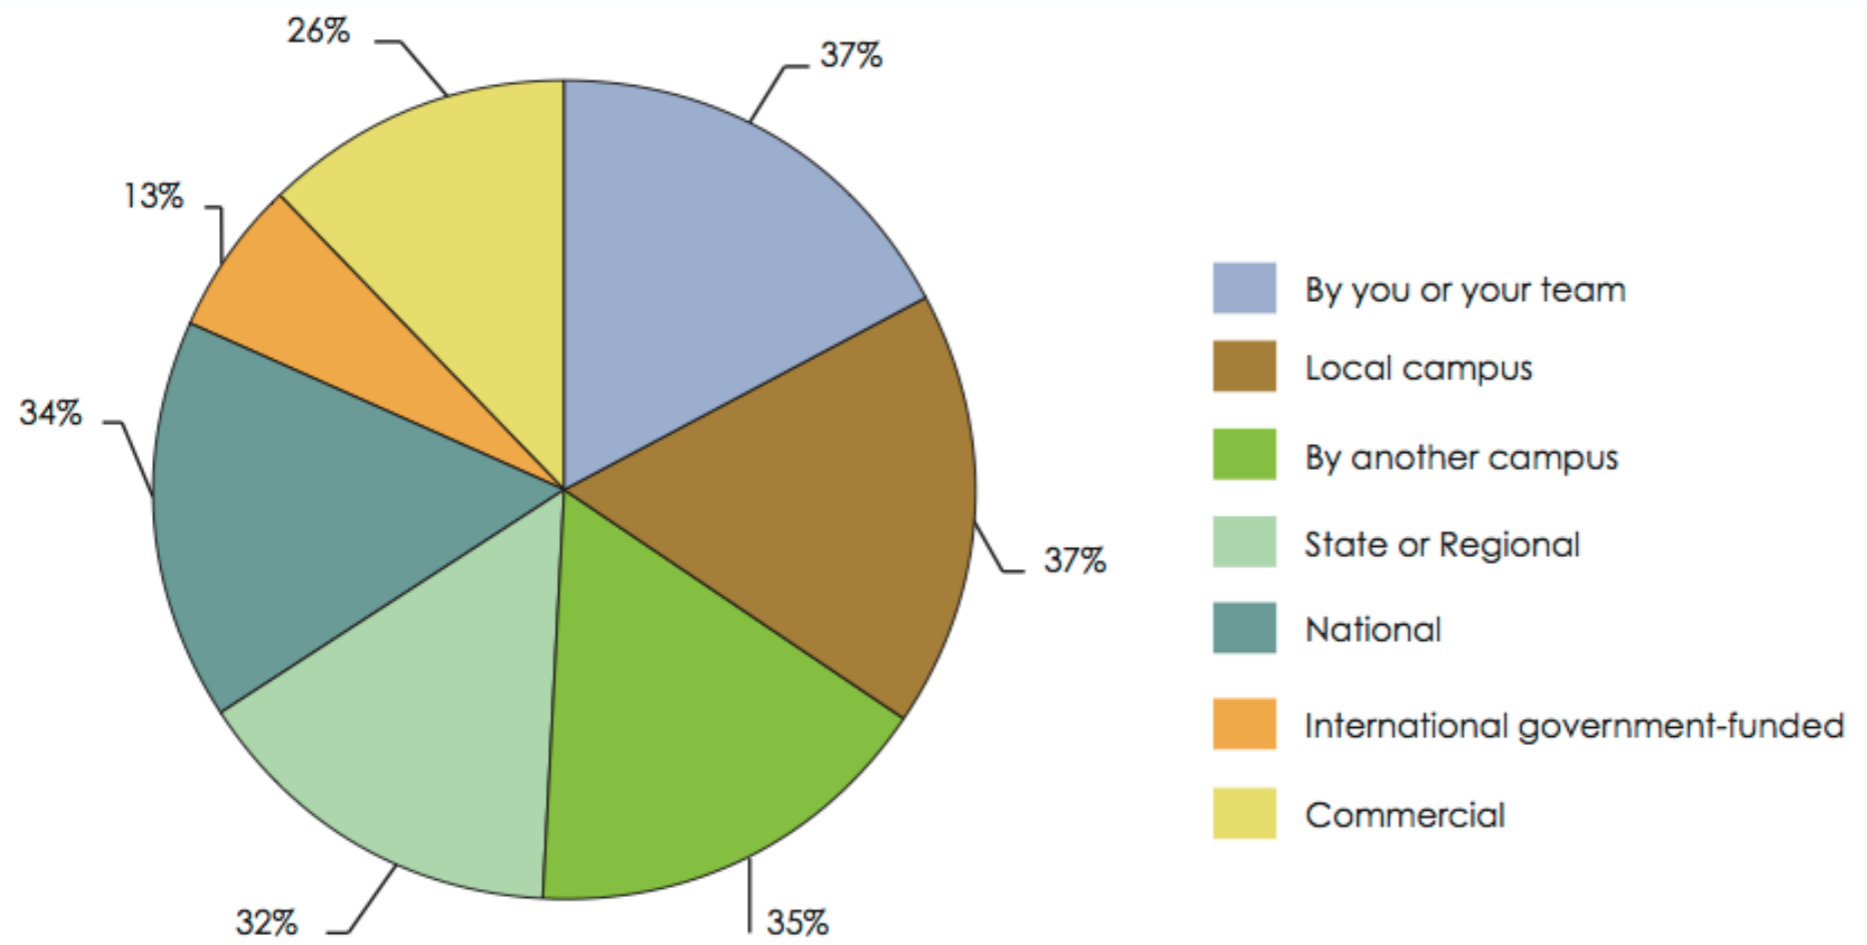
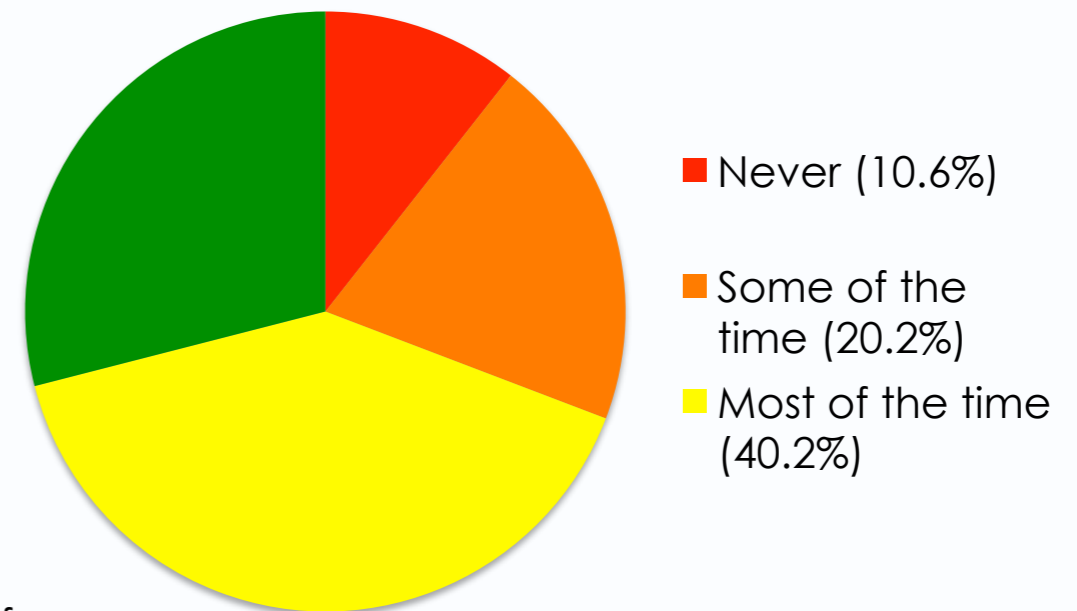


Figure and data from Campus Bridging: Software and Service Issues Workshop Report, McFee et al, <http://hdl.handle.net/2022/13070>

◆ Responses to asking if researchers had sufficient access to cyber infrastructure resources

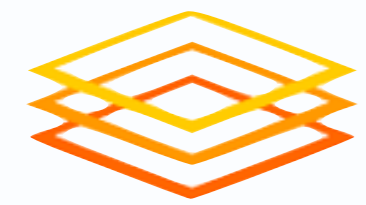


C.A. Stewart et al: Survey of cyberinfrastructure needs and interests of NSF-funded principal investigators. 2011. hdl.handle.net/2022/9917



The Gap

- ◆ “In our assessment, it is crucial to prevent the capability gap between these two extremes from growing. While “exascale” problems need to be solved for the large collaborations, it is equally necessary to ensure solutions are available for the many scientists challenged at the terascale and petascale on the passage to the exascale.”
- ◆ “To help our growing spectrum of users, we need to keep up with increasingly dynamic and heterogeneous environments while ensuring that domain scientists with limited computing expertise can use those environments.”

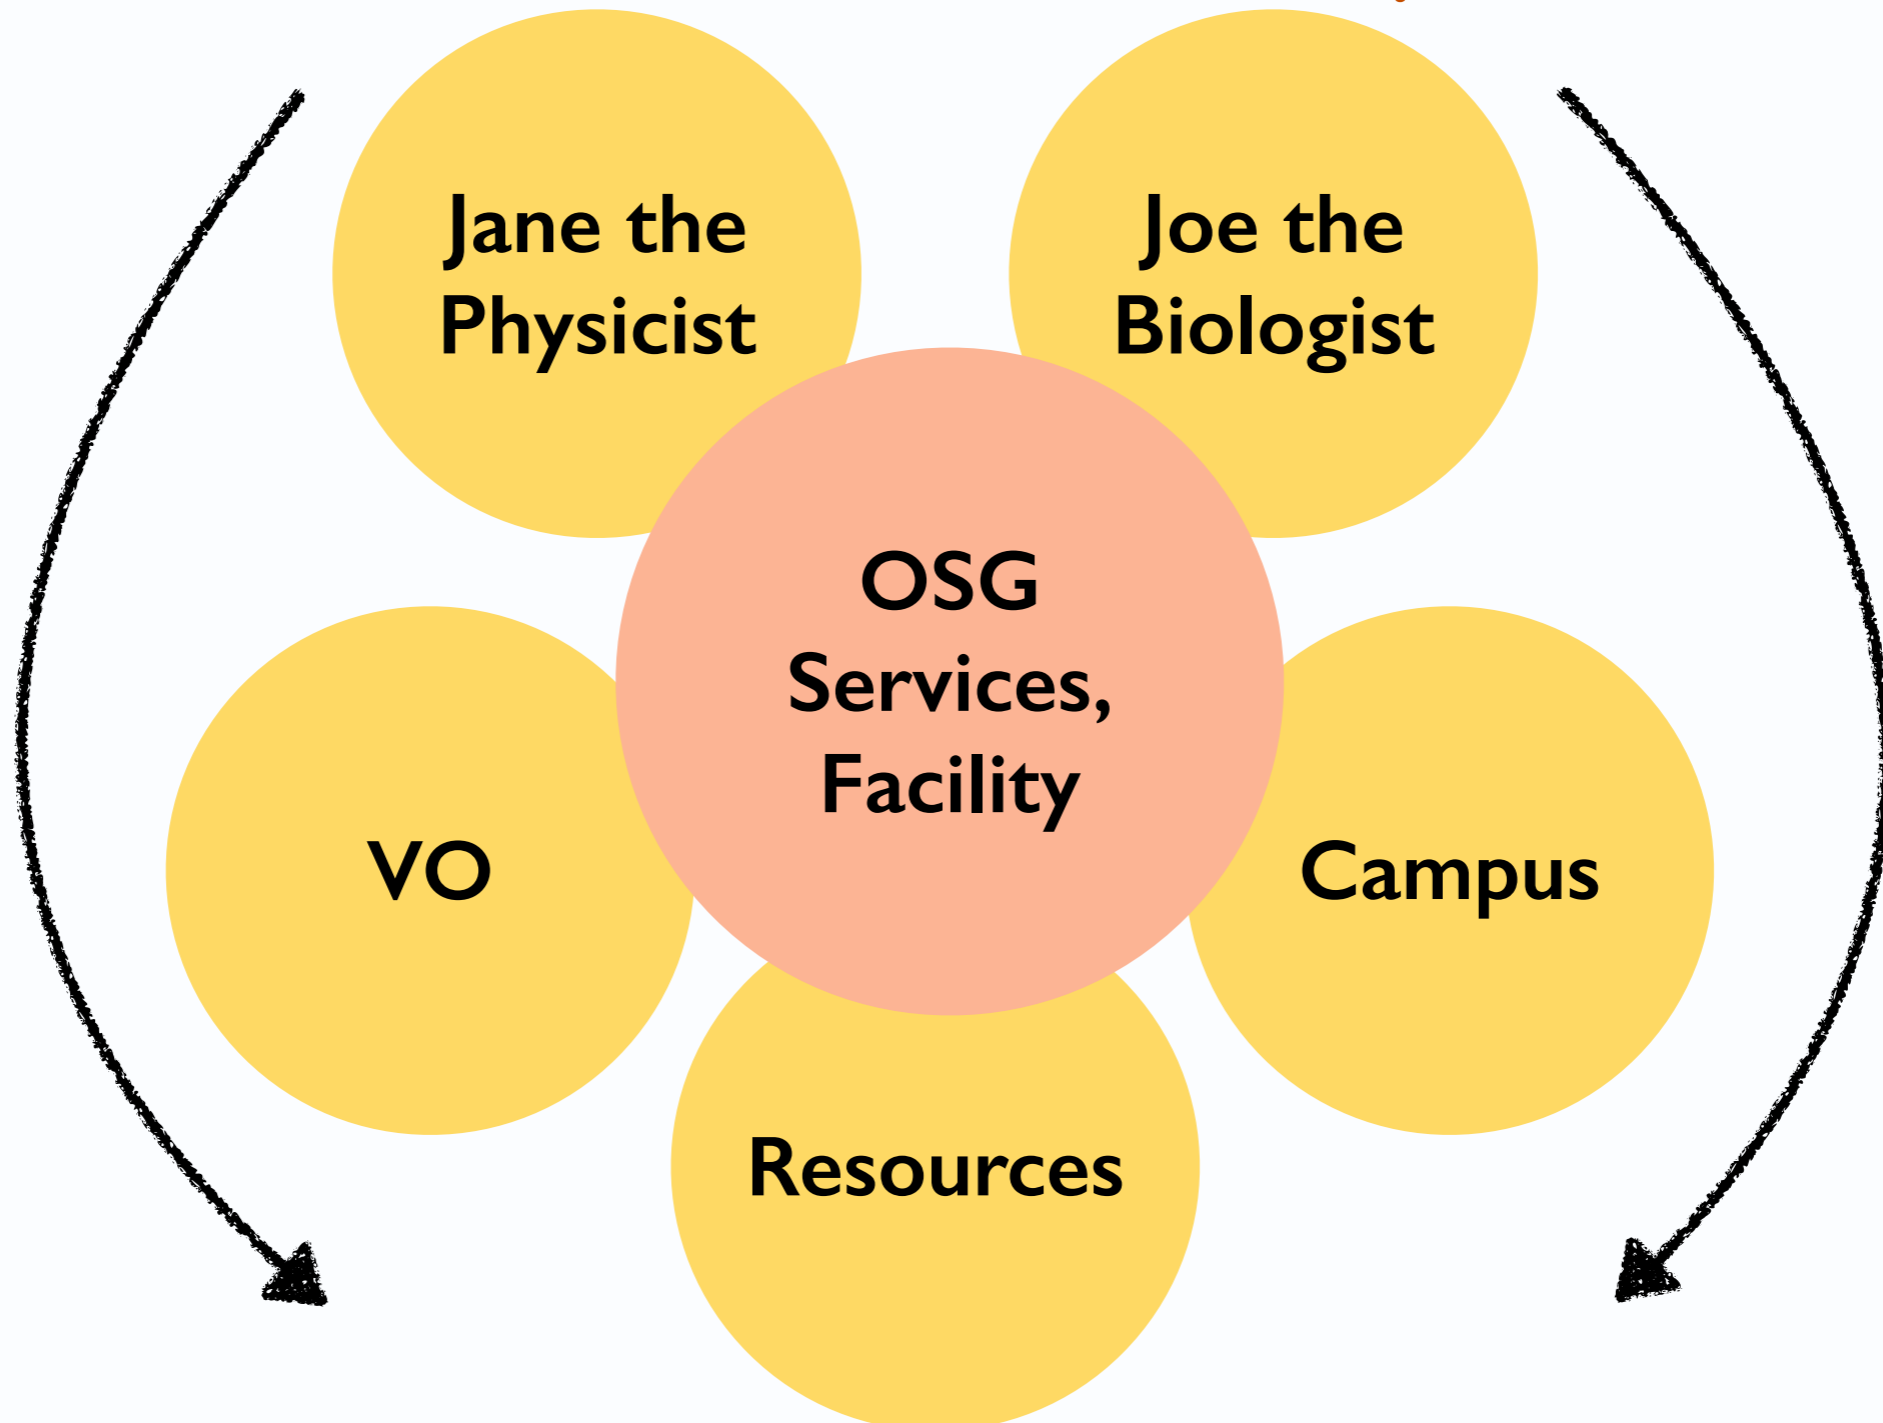


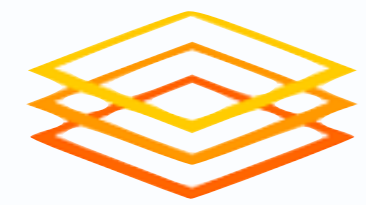
OSG Services and Facility Ecosystem

Open Science Grid

VO-centric Model

Campus User Model





“Mental Model” of OSG DHTC Facility

OSG “DHTC Facilities”

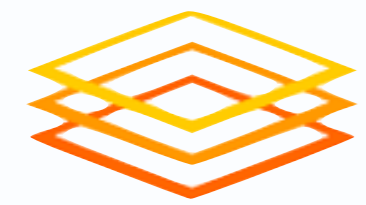
Tools and Services

VO Facilities:
Federate and Integrate
Enable Sharing

Open Facility:
Harvested
Resources

Provisioning Resources

Connection Layer
Network, Trust Relationships, Identities



“Mental Model” of OSG DHTC Facility

OSG “DHTC Facilities”

Tools and Services

VO Facilities:
Federate and Integrate
Enable Sharing

Open Facility:
Harvested
Resources

Provisioning Resources

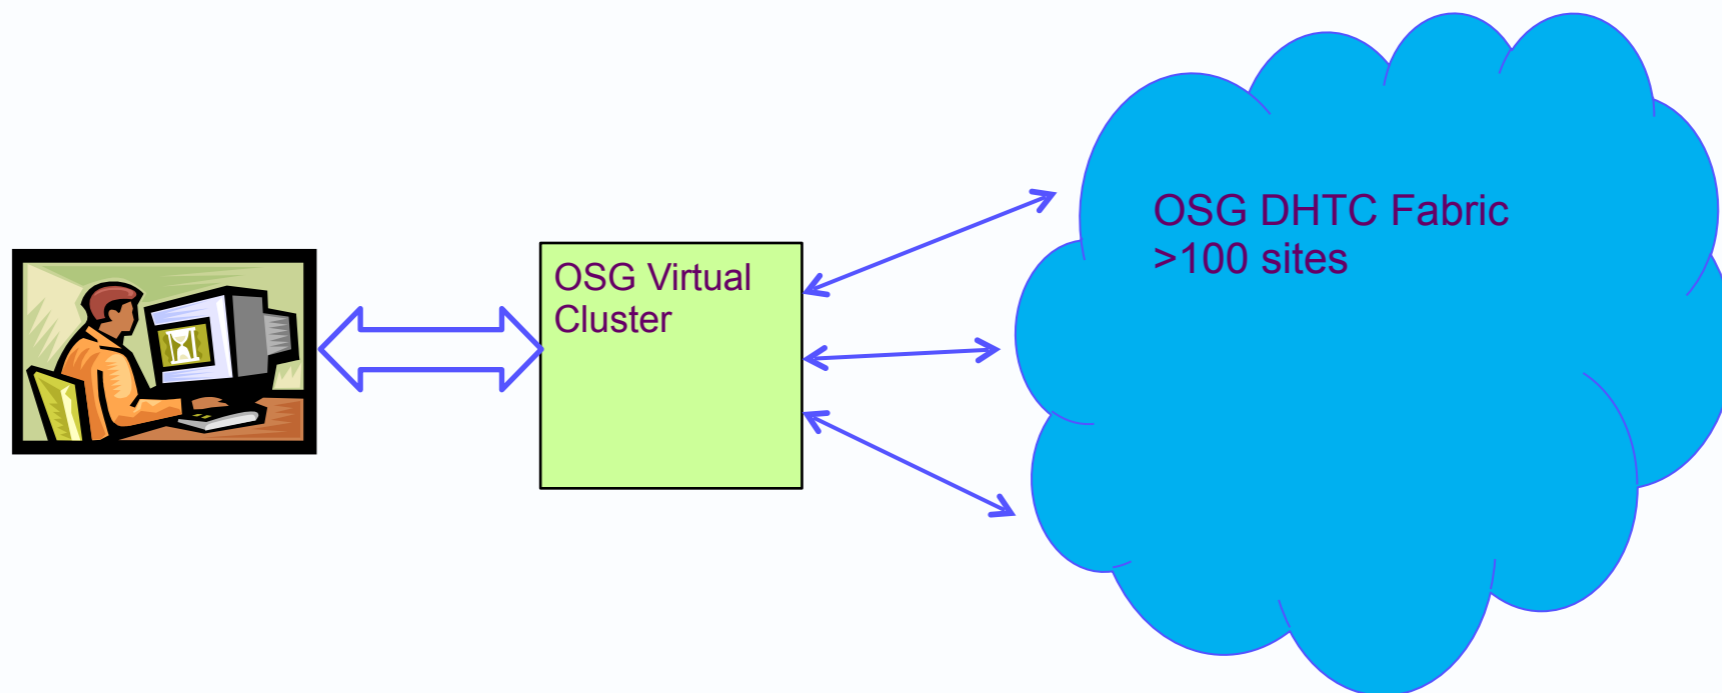
Connection Layer
Network, Trust Relationships, Identities



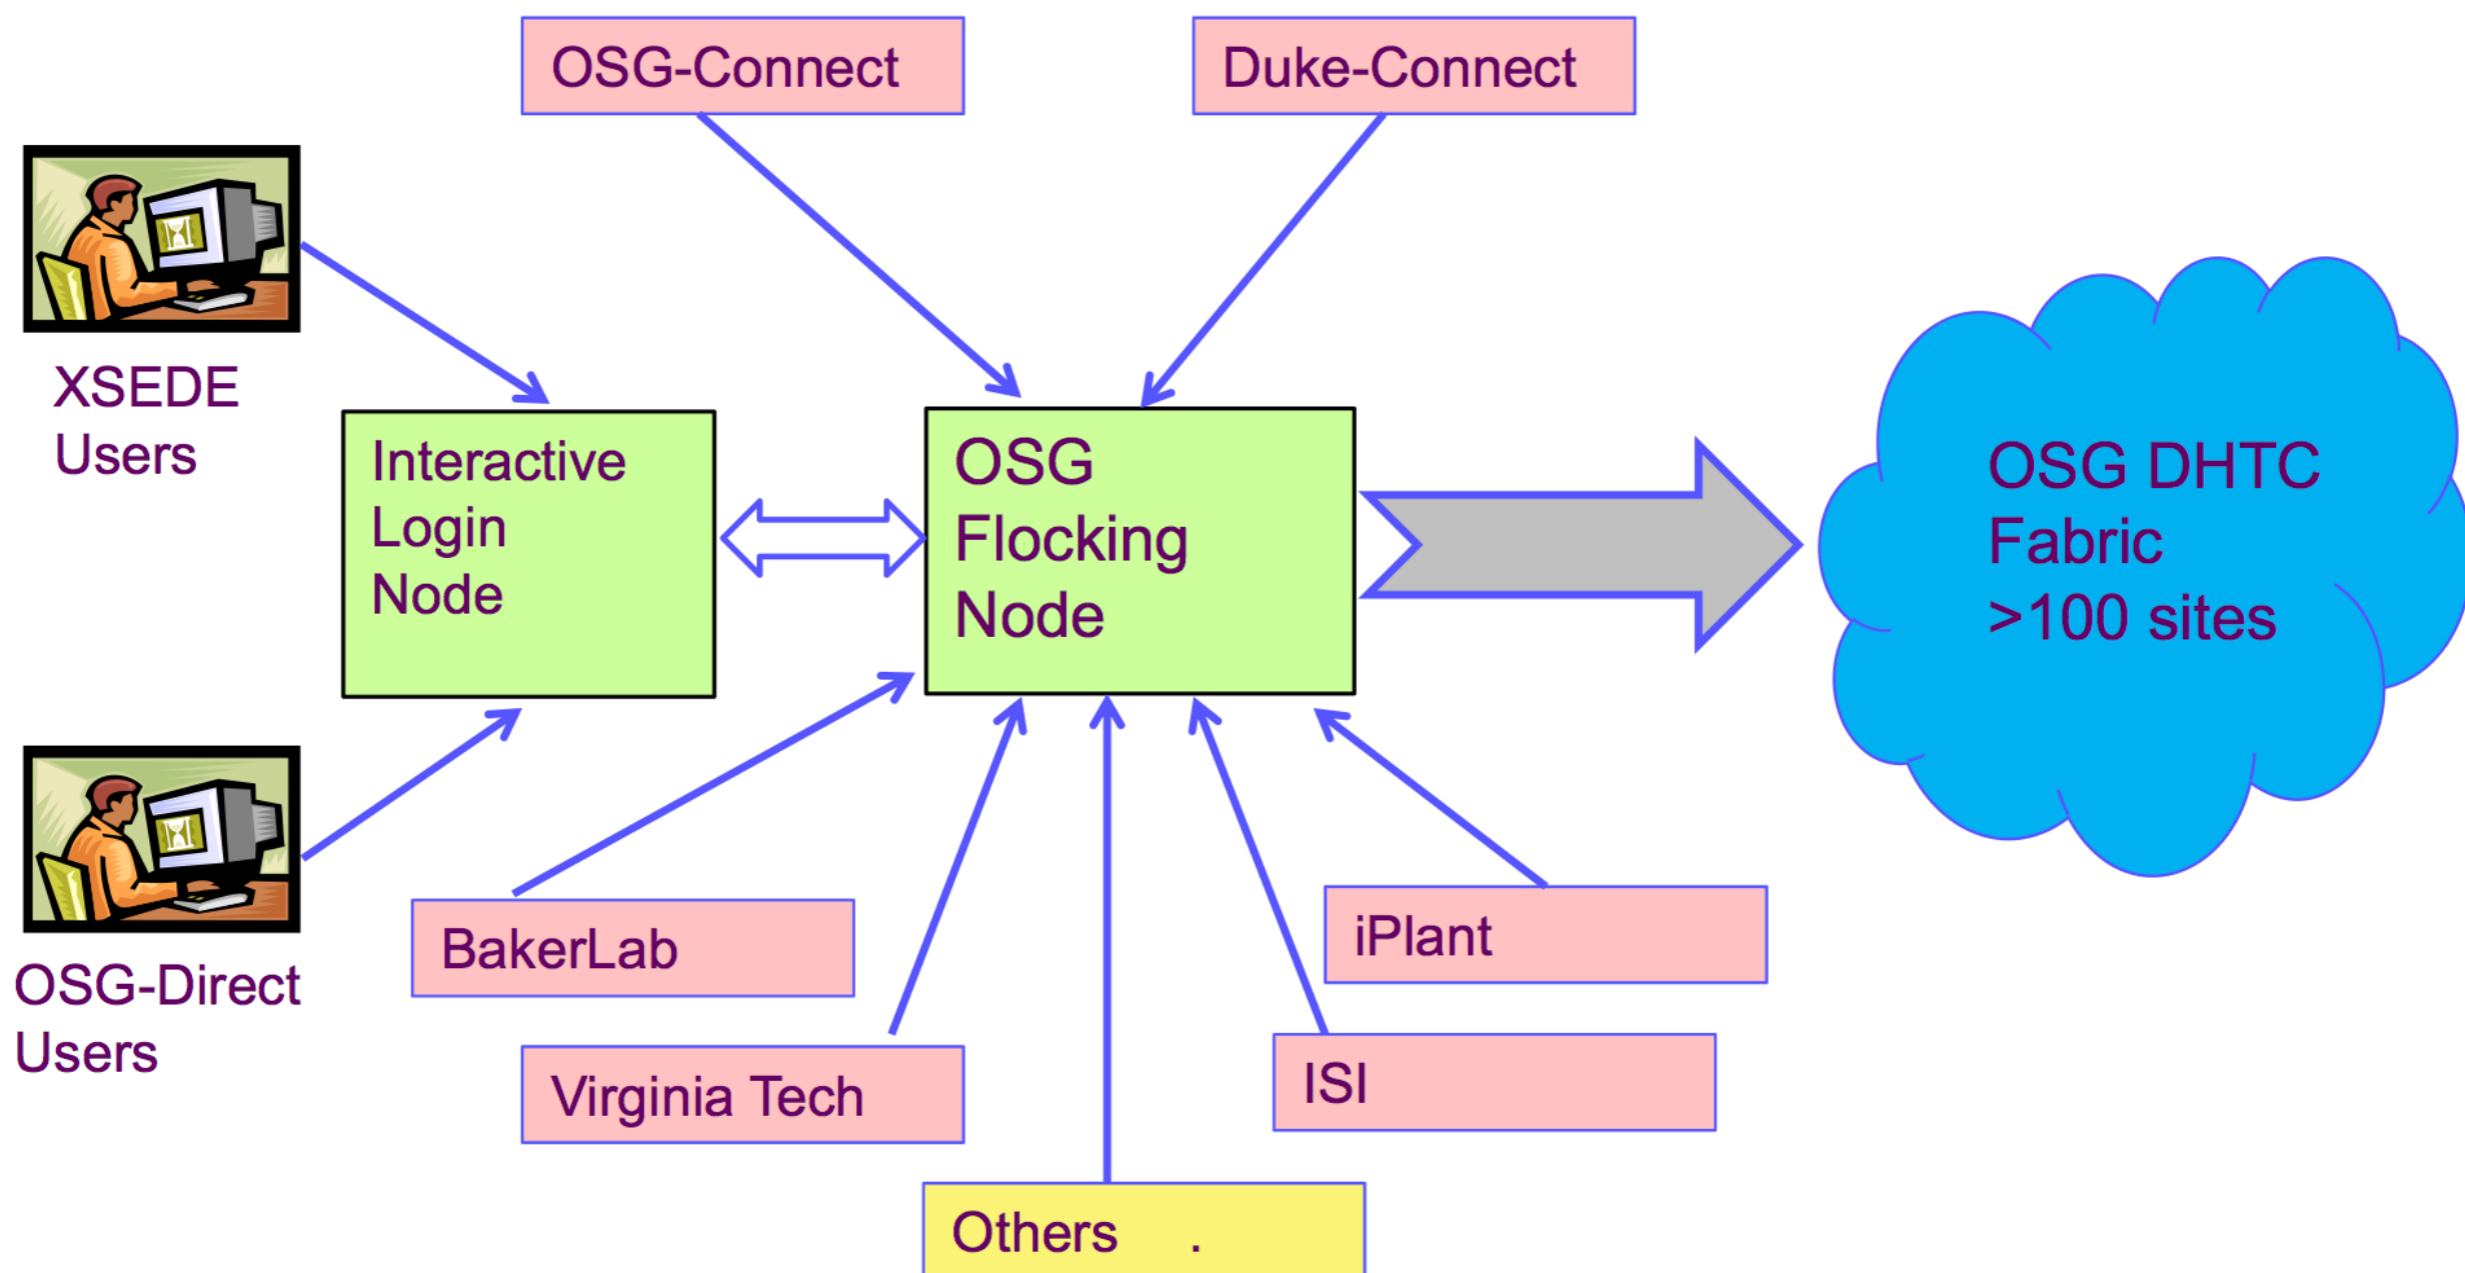
The OSG “Open Facility”

Open Science Grid

- ◆ A Distributed High-Throughput Computing Facility based on harvesting otherwise idle resources
 - ★ used by a quite large and diverse community of researchers and groups
 - ★ is being provided across the fabric, with resources harvested opportunistically
 - ★ the “OSG-VO” setup is now replacing engage and the focus of opportunistic cycles
 - ★ this facility is also the basis for OSG as an XD Service Provider



Current setup to allow access to OSG Fabric via OSG-VO



All access operates under the OSG VO using glideinWMS

Communities using those resources harvested from sites: Campus Users

◆ OSG VO:

★ XD Service Provider

- ◆ about 30M hours

★ OSG Connect users

★ a number of science applications

★ individual science groups, who also come in through the engage VO

◆ Users flowing into OSG from campus grids:

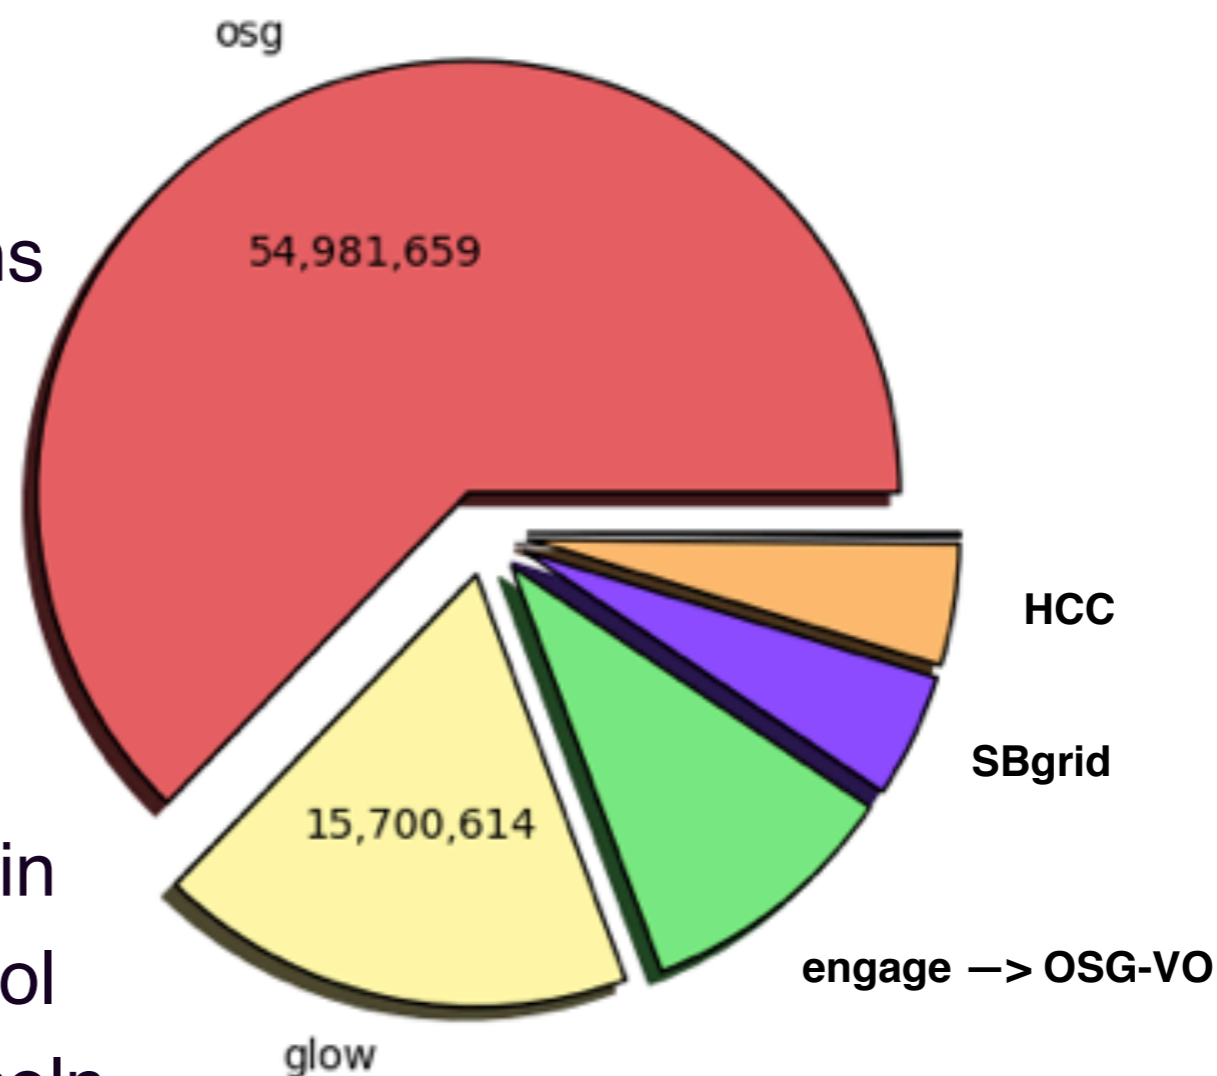
★ glow campus grid at U.Wisconsin

★ sbgrid at Harvard Medical School

★ hcc campus grid at Omaha/Lincoln

Wall Hours by VO (Sum: 87,652,040 Hours)

52 Weeks from Week 05 of 2013 to Week 05 of 2014



osg (54,981,660)
gluex (135,092)

glow (15,700,615)

engage (8,725,196)

sbgrid (4,103,054)

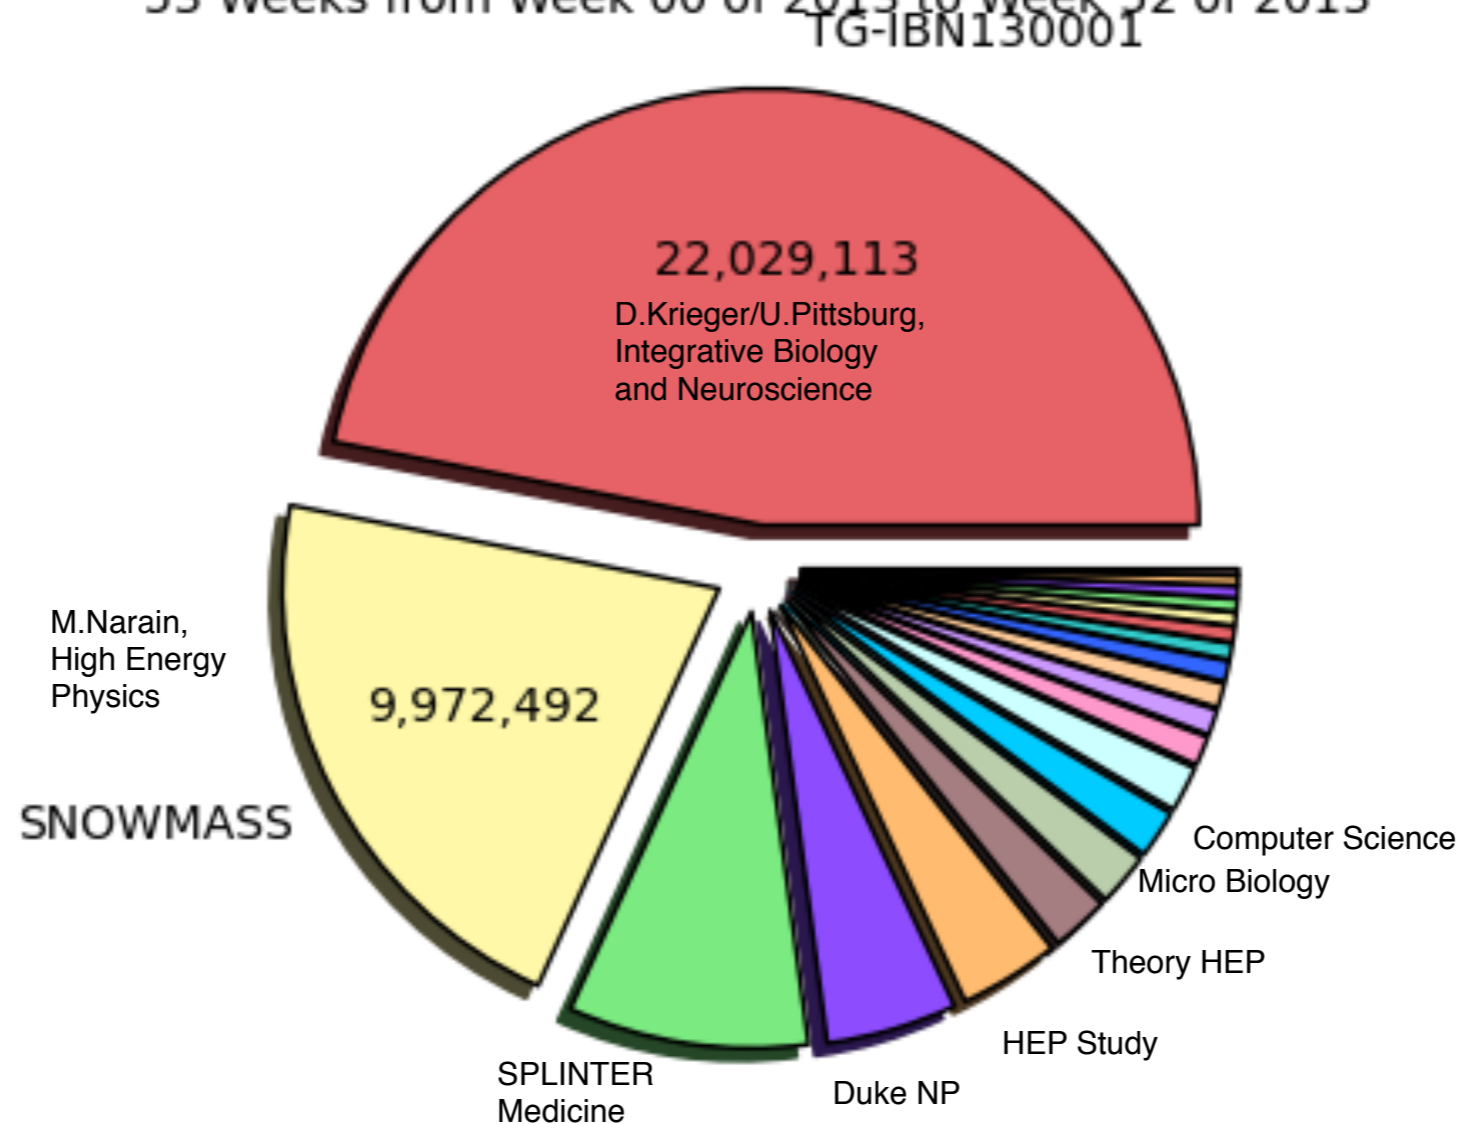
hcc (4,006,423)



Opportunistic Resourced: OSG-VO in 2013

Wall Hours by VO (Sum: 46,928,493 Hours)

53 Weeks from Week 00 of 2013 to Week 52 of 2013



TG-IBN130001 (22,029,113)
ECFA (1,744,646)
Other (797,936)
EIC (354,954)
TG-DMR130036 (212,059)

SNOWMASS (9,972,493)
TG-PHY110015 (1,004,429)
TG-PHY120014 (510,621)
UPRRP-MR (274,928)
IU-GALAXY (199,146)

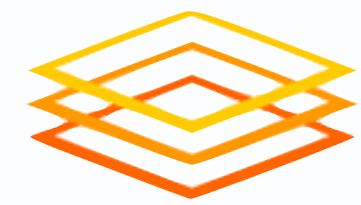
SPLINTER (4,174,511)
UMICH (925,155)
TG-TRA100004 (444,375)
TG-MCB100109 (262,308)
KNOWLEDGESYS (164,731)

DUKE-QGP (2,275,491)
RIT (836,827)
DETECTORDESIGN (421,086)
NESCENT (220,547)
DUKE (103,137)



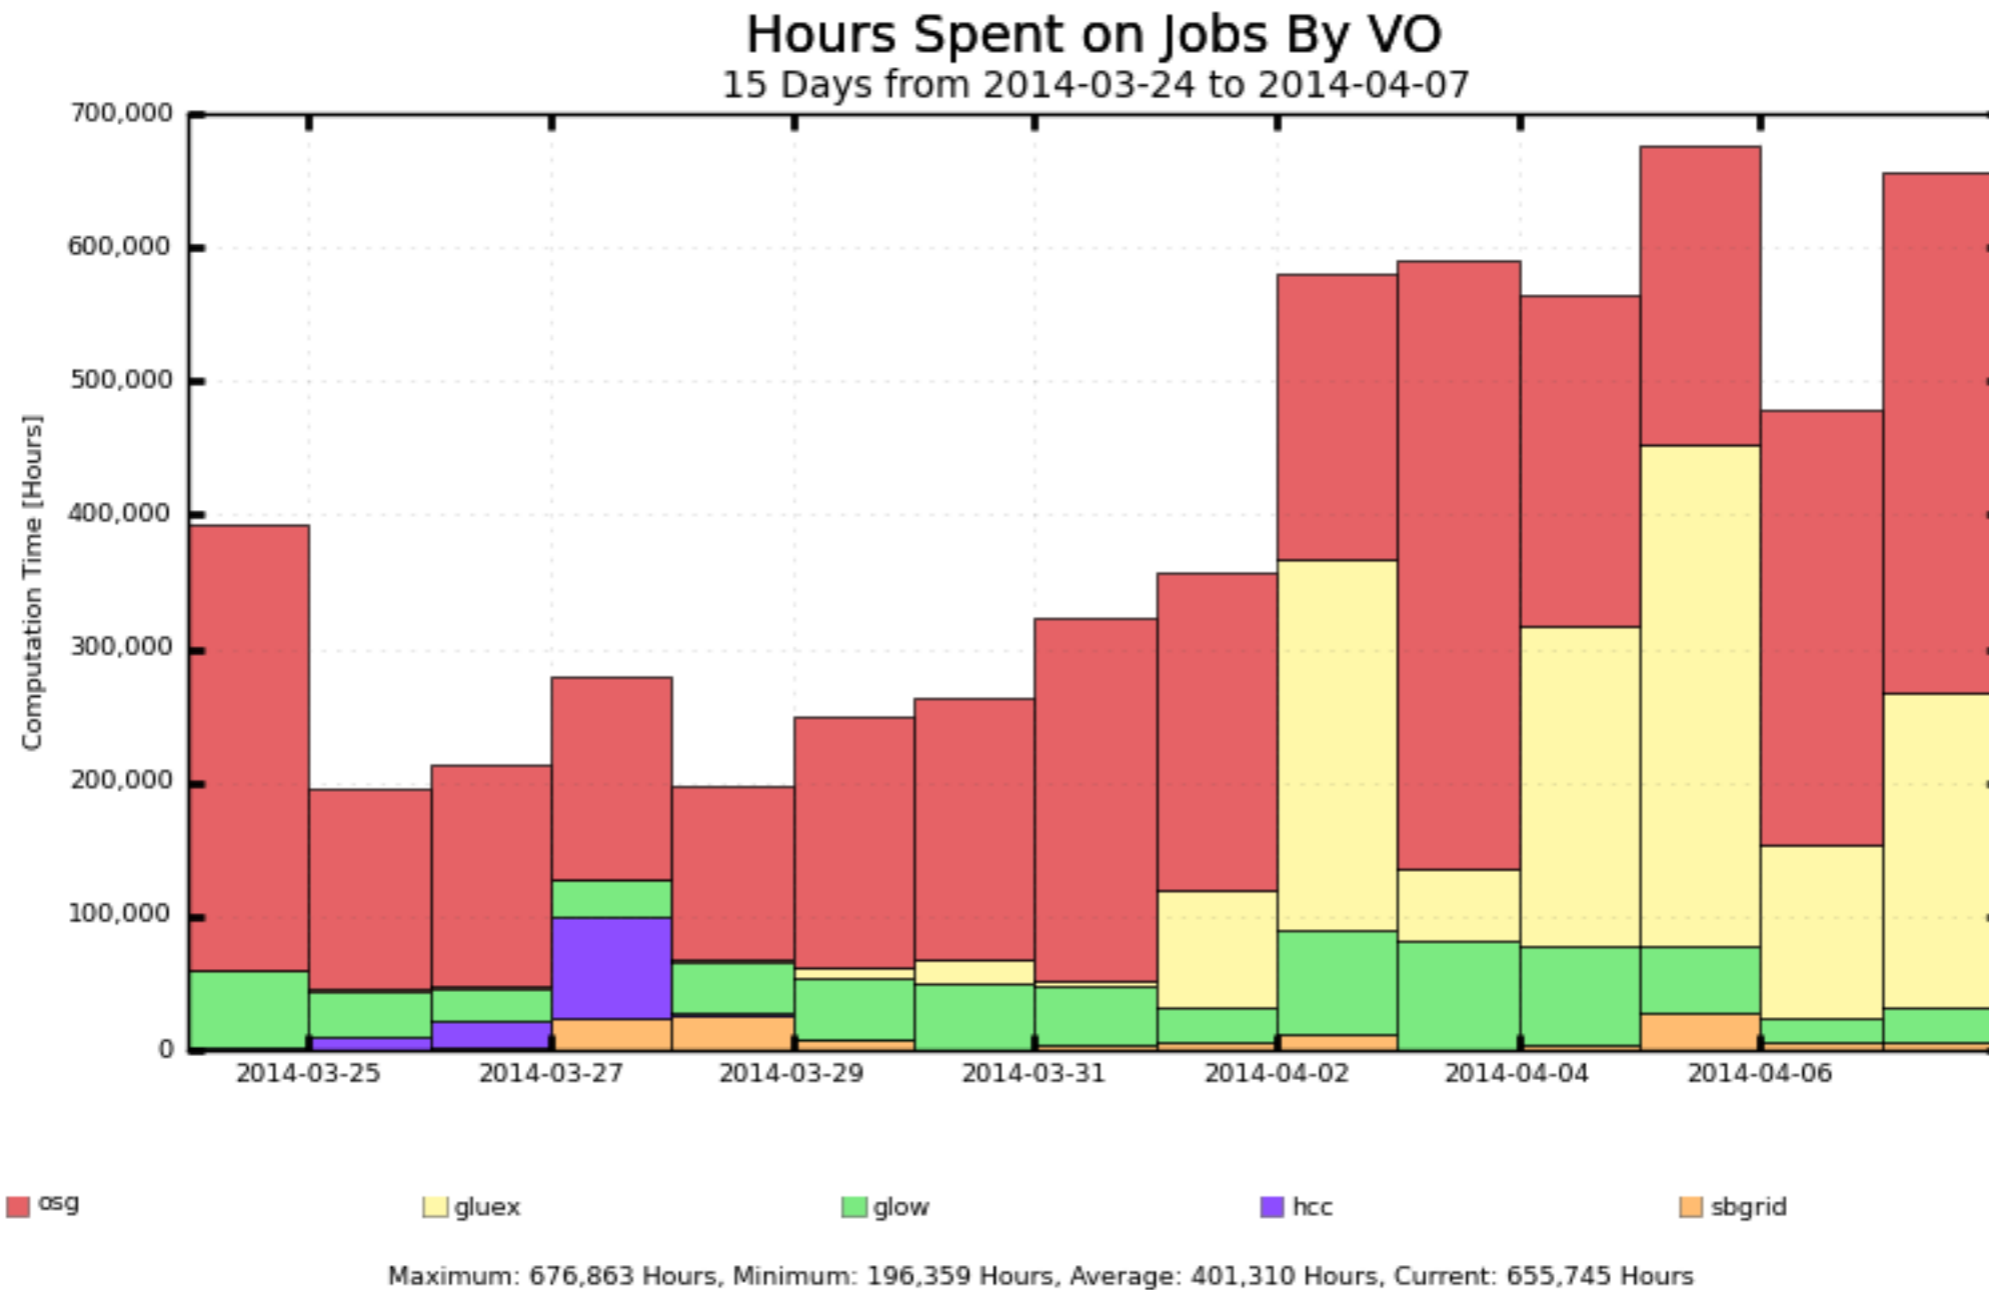
(from June/2013 staff retreat)

9
~~30M~~ hours!



Limit of the opportunistic use?

Open Science Grid



◆ not yet, it just doubled... (thx to Chander)

A number of open issues for OSG Open Facility

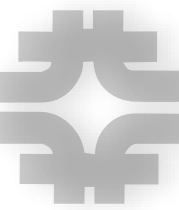
- ◆ Additional capabilities are needed to make this “production quality”, accountable, predictable, effective
 - ★ questions like, what actually is the limit of available resources?
 - ★ how do we prioritize, provision resources to users
 - ★ OSG-VO vs every VO harvesting on its own
 - ★ how do we add functionality, e.g. data services

- ◆ Add a wider range of customers, applications and use cases
 - ★ what additional services/capabilities are needed to make this an attractive facility
 - ★ e.g. should we spend effort to provide application “front-ends” like science gateways?
 - ★ e.g. what can/should we be doing regarding data access and data management?

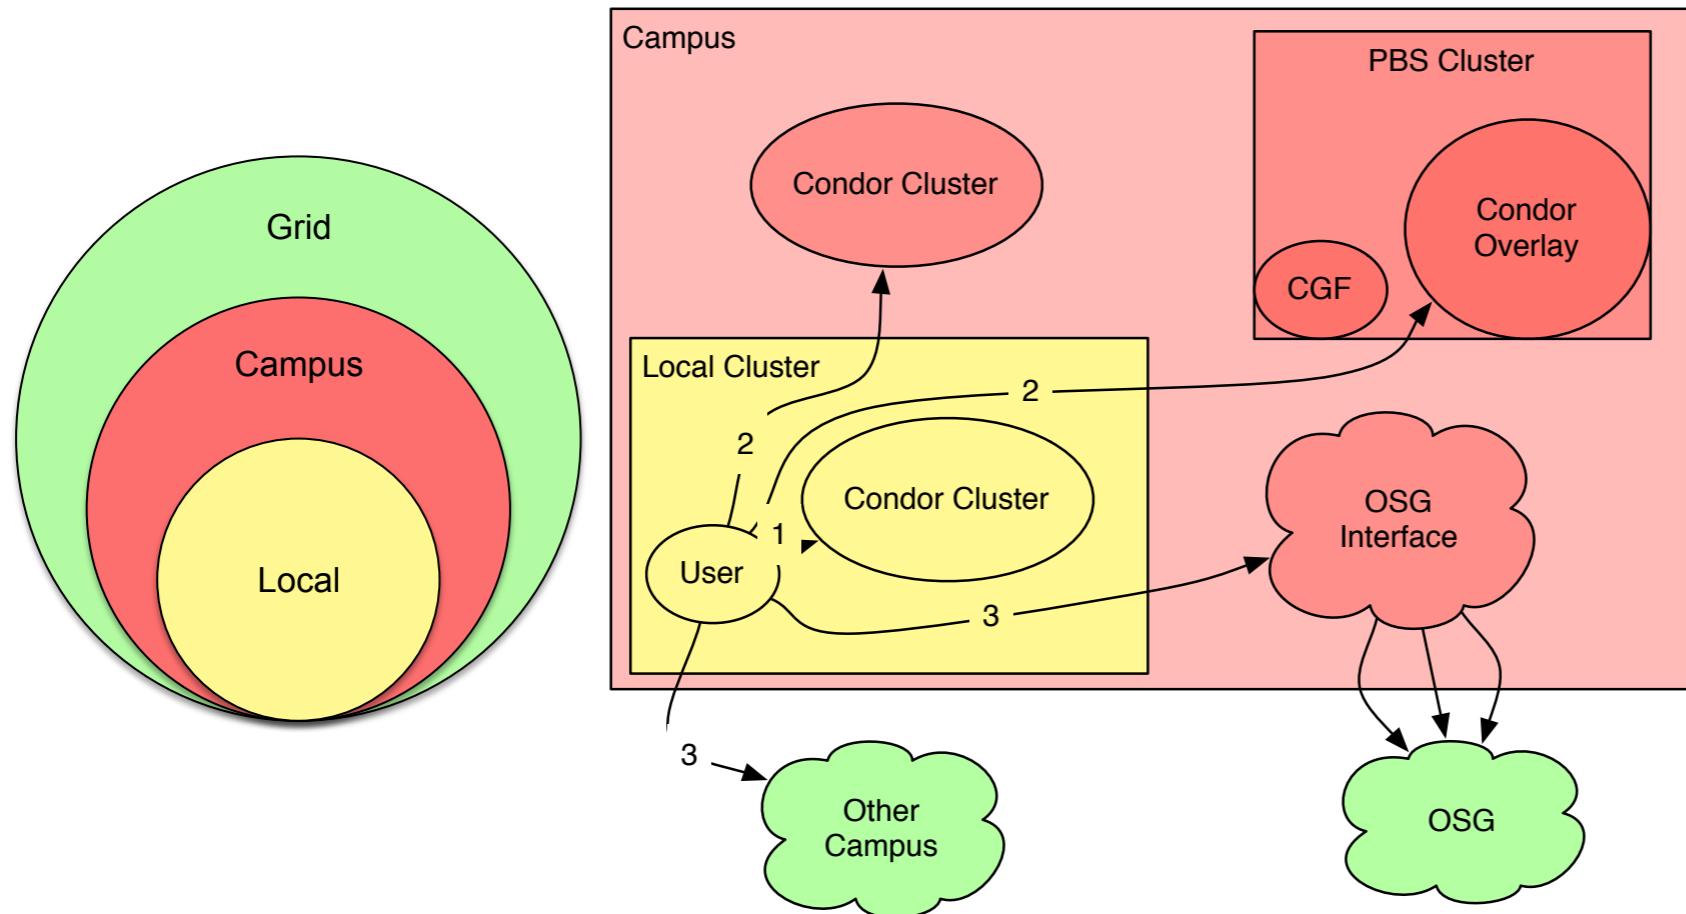
Developing OSG Strategies for Campus Users and Campus Grids

- ◆ Campus researchers club
 - ★ now discontinued, allowed individuals login-access to OSG Open Facility
- ◆ Campus Grids
 - ★ OSG services at sites to overflow from Campus into OSG Open Facility
- ◆ BOSCO “think (and submit) local, compute global”
 - ★ individual researcher access to OSG Open Facility
- ◆ XD service provider
 - ★ XD allocations to PIs/groups to use OSG Open Facility
- ◆ OSG Connect

Campus Grids in OSG, Example: Holland Computing Center (2013)



HCC Model for a Campus Grid



Me, my friends and everyone else

OSG AHM 2013



resources are combined

naha

submission protocol

to OSG sites across the
(cloud)

4 ensures high utilization

◆ D.Swanson, Director
Holland Computing Center

OSG AHM 2013





Open Science Grid

BOSCO to Enable Researchers to “Compute Globally”



BOSCO
Open Science Grid

BOSCO

About

News

Download

Installation

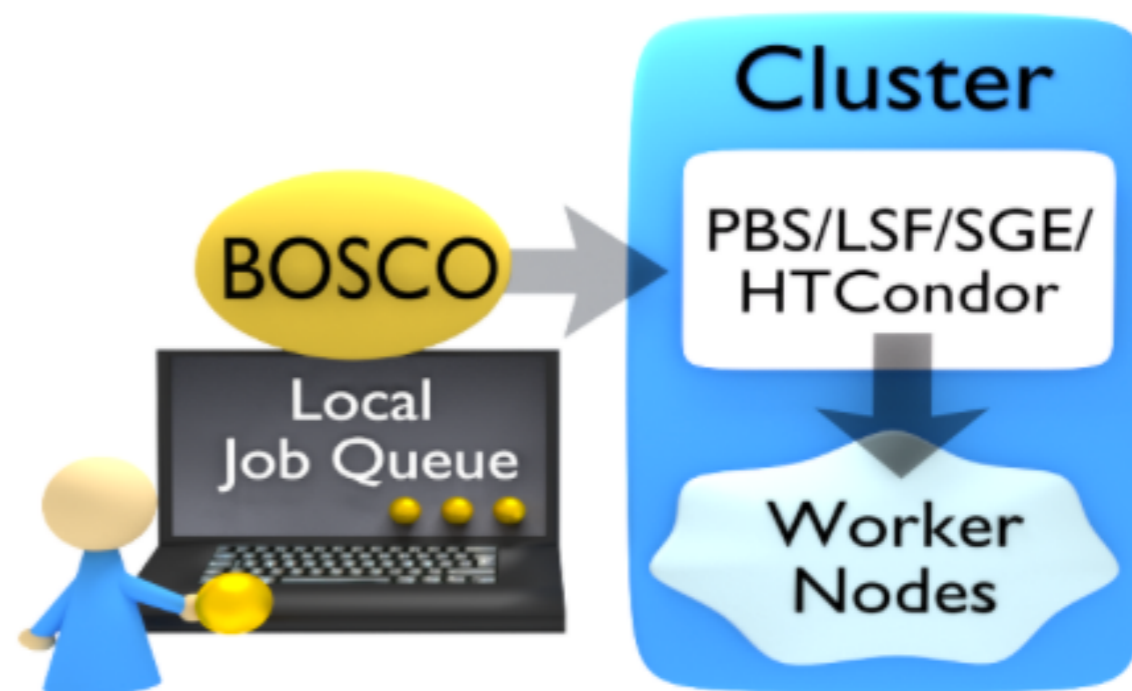
Contacts

BOSCO

Submit Locally. Compute Globally.

BOSCO makes it easy for you to access High Throughput Computing resources on your campus or the Cloud from your desktop.

[Download](#)



Workflow Management

Handles complex job dependencies

Multiple Clusters

Runs on LSF, PBS, SGE, and HTCondor

Simple Setup

Installed as non-root, simple config



BOSCO Successes and Failures

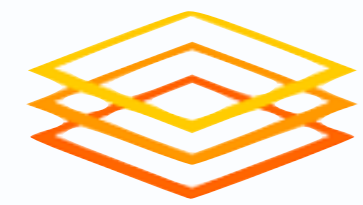
Open Science Grid

- ◆ Individual users at campuses can use the OSG through BOSCO
- ◆ A “Tool” approach:
 - ★ a software product that individual researchers can download and install
 - ★ to help users create larger, more inclusive Campus Grids, based on HTcondor
 - ★ goal of a downloadable self-installable and self-configuring robust tool to allow researchers HTC on whatever resources are easiest available to her/him — not yet fully achieved
- ◆ Great success as a power tool for campuses
 - ★ BOSCO is a very successful technology
 - ★ “HTcondor with ssh” —>moved back into HTcondor project
 - ★ widely used by experts
 - ◆ for accessing HPCs at SDSC, Stampede, NERSC
 - ◆ to enable OSG connect access to sites
 - ◆ many other uses



We're great at providing Services!

... and maybe not as great to provide products that users can pick up and run with



OSG Connect: Focus on Services

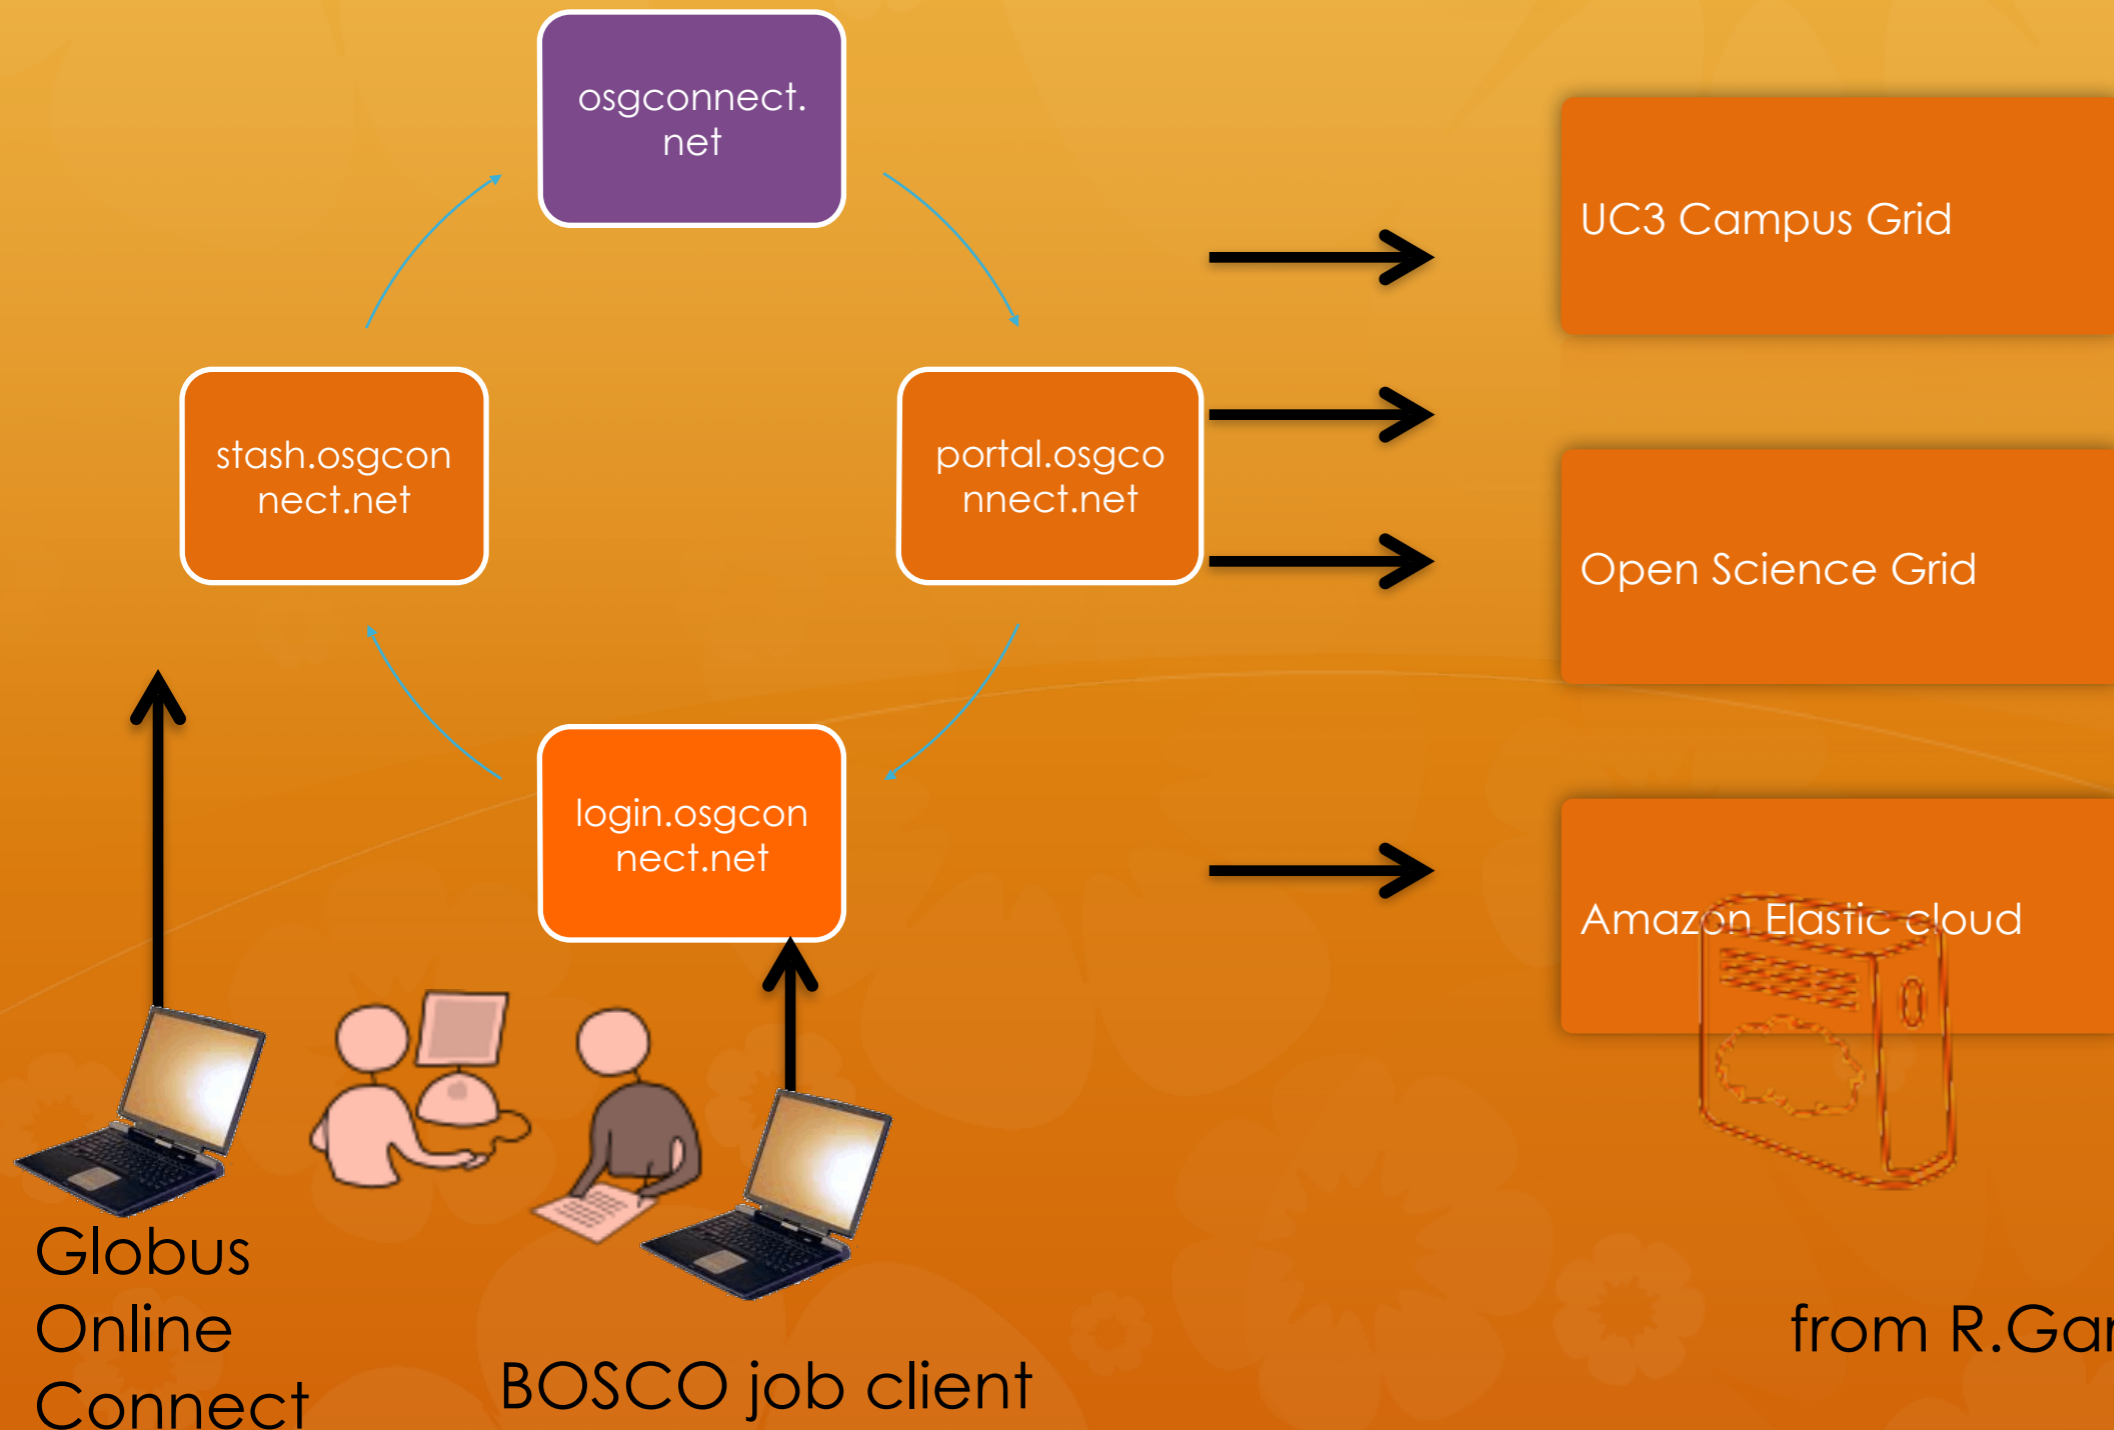
Open Science Grid

- ★ Core idea: provide a set of services that OSG runs for campuses and users
- ◆ OSG Connect “encapsulates” basic infrastructure services:
 - ★ from the network services (Globus online etc), to identity management services (bring you campus identity or your Globus ID, we’ll map it to your OSG identity)
 - ★ allow existing identities of researchers to be used across full set of services
 - ★ allow research communities to select services or tools that they require, without additional complexity
 - ★ setup easy access to a variety of facilities and resources
- ◆ Allows Campuses to create Campus Grids
 - ★ provide campuses with the services to connect to the OSG and a useful set of services that help setting up a campus grid or extending an application into a campus grid
 - ★ campuses don’t have to build their own — OSG is running services for sites
 - ★ initial set of services exists, two new campus grids being established
 - ★ getting some traction also in Atlas and CMS as a solution for Tier-3 etc

An OSG Connect Instance



osg connect



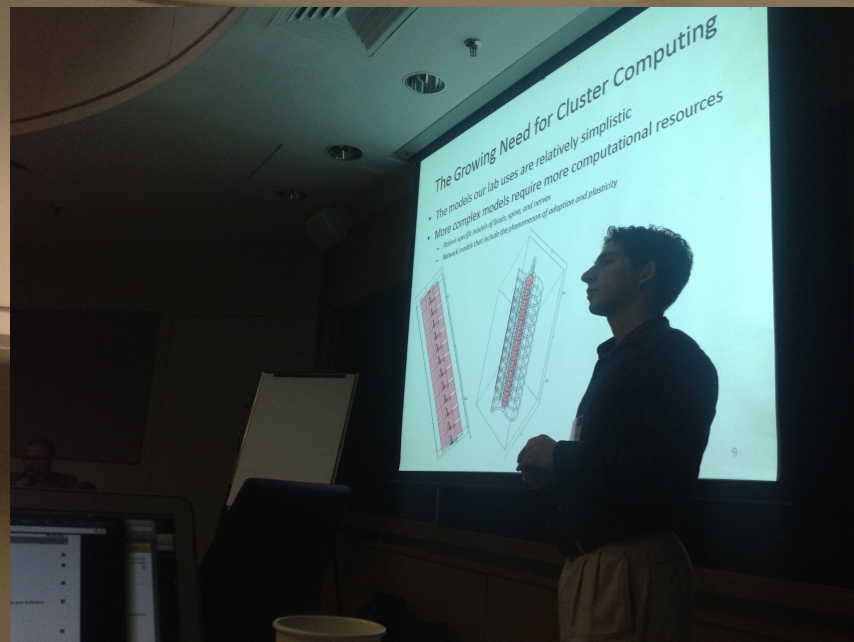
from R.Gardner/U.Chicago



OSG Connect as a Vehicle

Open Science Grid

- ◆ Extending the OSG eco system to the HPC community
 - ★ OSG is a XD service provider,
how about XD providing resources to OSG communities
 - ★ interface to XSEDE installations
 - ★ others should probably Piggy-back on Atlas Connect XSEDE initiative
 - ◆ glideinWMS and BOSCO access to DOE facilities
 - ★ HPC as a data producer
 - ◆ e.g. cosmology simulations
 - ◆ e.g. HEP simulation à la Stefan Höche (tomorrow)
- ◆ To establish more Atlas (and hopefully CMS) Tier-3 centers
 - ★ set up the list of services on “CI Connect” installation
 - ★ then connect to campuses and resources
 - ★ Is this also a model for future “VO facilities”, that bring in new sites?
- ◆ OSG Connect could also be a vehicle to establish new service offerings
 - ★ Panda service -- next step for OSG to help with packaging
 - ★ then deploy service for (non-Atlas) Panda users



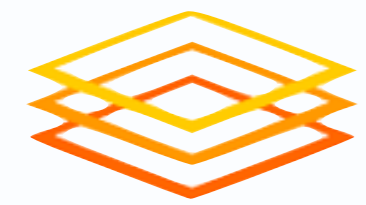
**Campus Infrastructure Community
Workshops like the one at Duke
brought a number of new users
in it's wake**



How About Data?

Open Science Grid

- ◆ Large VOs offer users sophisticated data infrastructures
 - ★ dynamic data placements and high throughput robust data transfers
 - ★ data discovery and metadata tools
- ◆ OSG has not yet succeeded (nor tried) to make such available to “long tail”
- ◆ However, the real killer feature: Connectivity and Network Throughput
 - ★ including networks and distributed IDM
 - ★ enabling remote data access, federating storage systems and allowing global access to locally managed data
- ◆ We should start with something that's based on local data and that's simple
 - ★ OSG Connect already has a “stash” service
 - ★ Should add a data archiving service, e.g. at Fermilab
- ◆ These locally provided data services could becoming imminently useful through connectivity and federation (and possibly squid-based caching?)
 - ★ a simple general data archival storage facility could have a big impact
 - ★ sustainable, dependable, accessible archive, offer open access to the data
 - ★ a possible focus for data preservation activities, long-term guardianship of data

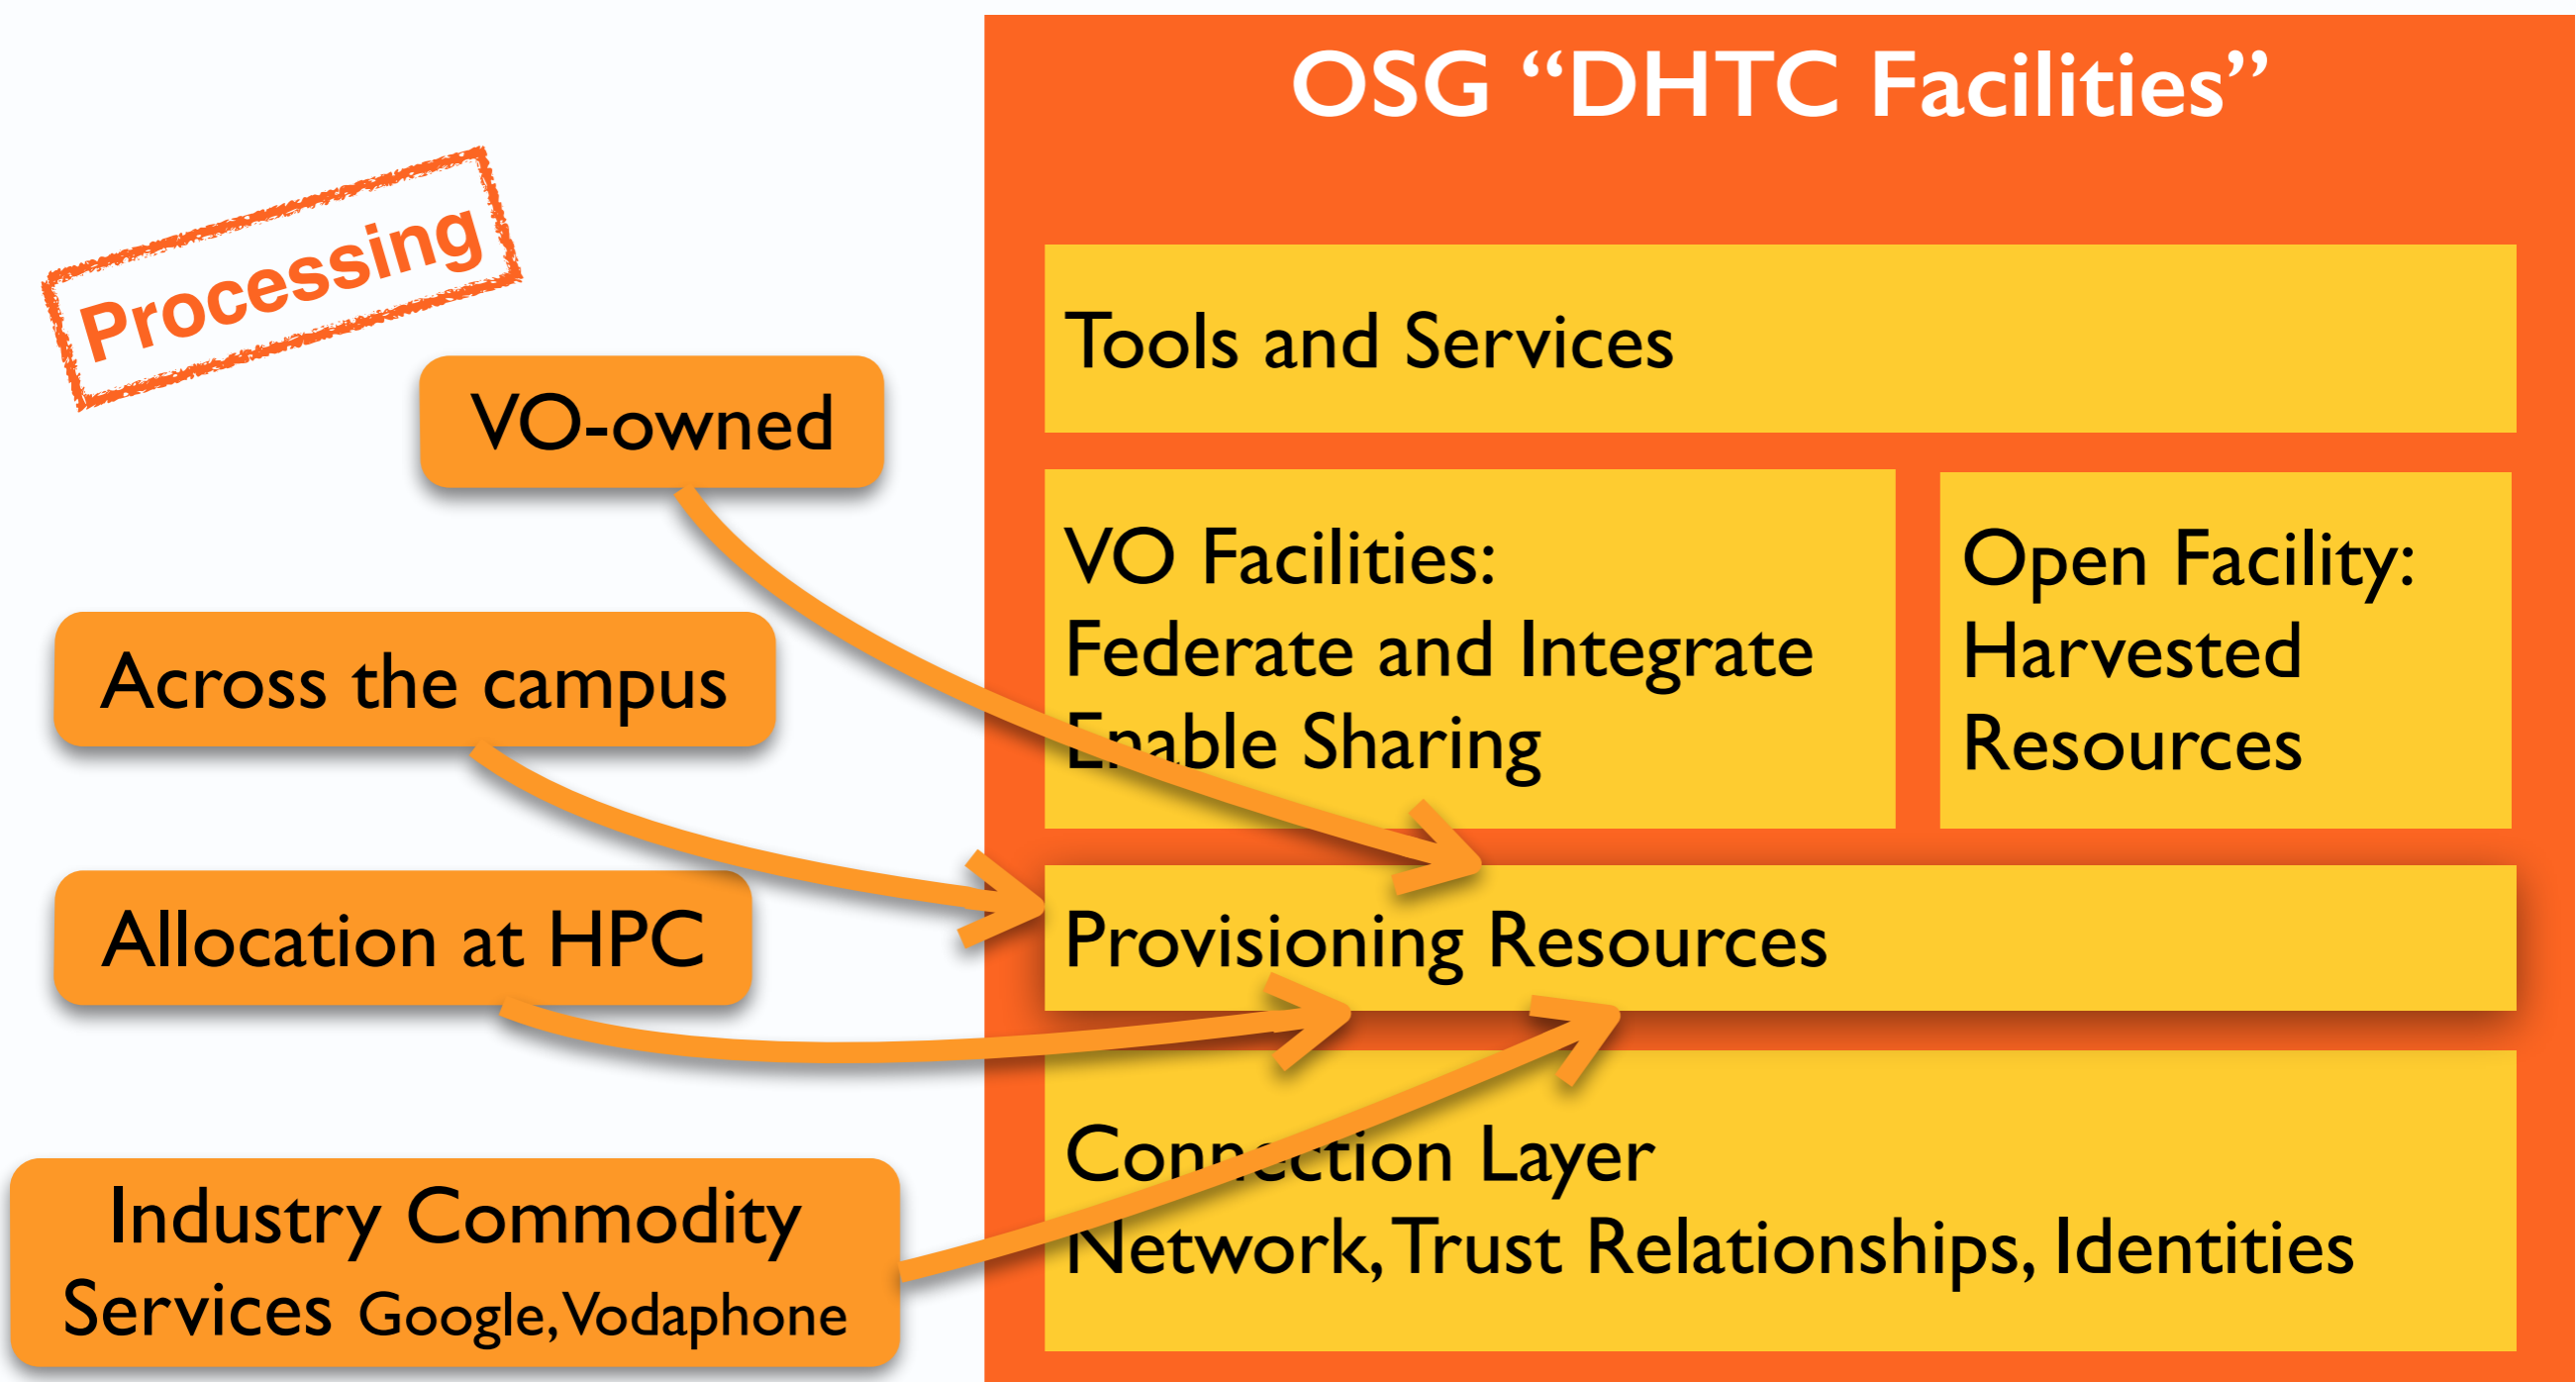


“Mental Model” of OSG DHTC Facility





Bringing in New Types of Resources



◆ Focus on *Resource Provisioning*

- ★ Statically federated resources need to be integrated with dynamically allocated resources causing new challenges for resource planning, acquisition, provisioning



Resource Provisioning to Campus Users

- ★ Can OSG assist VOs provision resources with different costs?
 - ◆ clouds are a “business model” that Resource Providers use to provide resource
 - ◆ for user, “costs” for using a specific compute resources are not necessarily money
 - ◆ other examples are CPU power, network locality, allocation
- ★ Should OSG act as an intermediary, between VOs and resources?
 - ◆ in this model, instead of OSG going to NERSC to ask for an allocation for it's users, e.g. Atlas has the allocation and OSG helps to manage and provision it to the Atlas workflows
- ★ OSG does not yet have a flexible “provisioning” capability
 - ◆ although it has been discussed as a natural extension addressing a number of shortcomings
 - ◆ raises new questions of liability, accountability, audit-ability for the “money” spent?
 - ◆ can this model be extended to “enabling acquisition” of commercial cycles o behalf of an OSG member
 - ◆ maybe even be extended to the provisioning of opportunistic resources



Success of OSG Campus Program is about minimizing obstacles for users

- ◆ Rob G: once users adapt they don't look back
 - ◆ So we have to be at our best game to make users adapt
 - ◆ OSG Connect, OSG-VO and and seamless access to the overlay job exec environments, XSEDE-SP, remote data access allowing local data management, going away from user certs, etc
- all are great ways to minimize impedance and get campus users to take the gateway drug and get high with their science on DHTC
- ◆ Thanks to a great team moving OSG forward towards the next 3 years!